DETECTION OPTIMIZATION OF THE PROGRESSIVE MULTI-CHANNEL CORRELATION ALGORITHM USED IN INFRASOUND NUCLEAR TREATY MONITORING

THESIS

Anthony M. Runco, Second Lieutenant, USAF

AFIT-ENG-13-M-42

AFIT-ENG-13-M-42

DETECTION OPTIMIZATION OF THE PROGRESSIVE MULTI-CHANNEL CORRELATION ALGORITHM USED IN INFRASOUND NUCLEAR TREATY MONITORING

THESIS

Presented to the Faculty

Department of Electrical Engineering

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the

Degree of Master of Science in Electrical Engineering

Anthony M. Runco, BS

Second Lieutenant, USAF

March 2013

AFIT-ENG-13-M-42

DETECTION OPTIMIZATION OF THE PROGRESSIVE MULTI-CHANNEL CORRELATION ALGORITHM USED IN INFRASOUND NUCLEAR TREATY MONITORING

Anthony M. Runco, BS
Second Lieutenant, USAF

Approved:

_____   _____
James A. Louthain, Lt Col, USAF (Chairman)   Date

_____   _____
Michael A. Temple, PhD (Member)   Date

_____   _____
Dean A. Clauter, PhD (Member)   Date

AFIT-ENG-13-M-42

Abstract

This thesis develops methods to determine optimum detection thresholds for the Progressive Multi-Channel Correlation (PMCC) algorithm used by the International Data Centre (IDC) to perform infrasound station-level nuclear-event detection. Receiver Operating Characteristic (ROC) curve analysis is used with real ground truth data to determine the trade-off between the probability of detection (P_D) and the false alarm rate (FAR) at various consistency detection thresholds. Further, statistical detection theory via maximum a posteriori and Bayes cost approaches is used to determine station-level optimum "family" size thresholds of grouped detection "pixels" with similar signal attributes (i.e. trace velocity, azimuth, time of arrival, and frequency content) before the detection should be considered for network-level processing. Optimum family sizes are determined based upon the consistency threshold and filter configuration used to filter sensor data prior to running the detection algorithm. Finally, this research generates synthetic signals for particular array configurations, adjusts the signal-to-noise ratio (SNR) to determine the SNR failure levels for the PMCC detection algorithm, and compares this performance to the performance of fielded infrasound stations with similar configurations. For the fielded stations studied, PMCC was able to detect signals with post-filtered SNRs as low as 2 dB, which represented approximately 2 dB better (lower) performance than as indicated by the synthetic test results.

iv

To the usual suspects, my parents, and rightfully so… To my father, who retired from a life of service and remains unconditionally invested in my and my brother's futures – no matter how old we get. To my mother, who will always be the most compassionate, caring, and altruistic person I know. No one can embark on the road of life alone, and, without you both, I would not have even had the opportunity to write this thesis. Without you both, I would surely be stuck in a roadside ditch.

Acknowledgments

Very simply, I am indebted to my advisor, Lt Col James Louthain. He remained calm, cool, and collected in what initially seemed like a daunting task in an unfamiliar field for both of us. His quiet confidence in me encouraged me throughout this research process in more ways than he knows. His insights during our many thought-provoking conversations proved invaluable, and his approachability should serve as a model to advisors everywhere. Although he outranks me in title and in experience, he always respected my opinions and made me feel as if I was one of his academic peers. For this, I am grateful. In the military, we respect the *rank* of a superior officer. In Lt Col Louthain's case, I also respect the *man*.

Further, I would be remiss not to express my earnest gratitude to Dr. Clauter, who worked tirelessly to launch this research beyond the ground-level planning stages. In this regard, I must also thank Dr. Arrowsmith at Los Alamos National Laboratory for providing me with ground truth data he collected using his own impressive detection program, InfraMonitor. Last, but not least, I would like to thank Dr. Temple for serving on my thesis committee and for offering advice that ultimately improved the quality of this written document.

Anthony M. Runco, Jr.

Table of Contents

List of Figures

x

List of Tables

Acronyms and Abbreviations

AFTAC	Air Force Technical Applications Center
AWGN	Additive white Gaussian noise
CTBT	Comprehensive Nuclear-Test-Ban Treaty
CTBTO	Comprehensive Nuclear-Test-Ban Treaty Organization
DMZ	Demilitarized Zone
FAR	False alarm rate
FDR	False discovery rate
FK	Frequency-wavenumber
GT	Ground Truth
IDC	International Data Centre
IMS	International Monitoring System
LANL	Los Alamos National Laboratory
LRT	Likelihood ratio test
LTA	Long-term-power-average
MAP	Maximum a posteriori
MUSIC	Multiple-Signal Characteristic
PMCC	Progressive Multi-Channel Correlation
PSD	Power spectral density
ROC	Receiver Operating Characteristic
SEL	Standard event list
SNR	Signal-to-noise ratio
SOI	Signal-of-interest
STA	Short-term-power-average
VPN	Virtual private network

DETECTION OPTIMIZATION OF THE PROGRESSIVE MULTI-CHANNEL CORRELATION ALGORITHM USED IN INFRASOUND NUCLEAR TREATY MONITORING

I. Introduction

The Air Force Technical Applications Center (AFTAC) has long used networks of nuclear event detection sensors to detect nuclear tests carried out anywhere on the globe. In its mission to achieve information superiority, AFTAC has historically combined data garnered from seismic and infrasound networks to improve location estimates for nuclear events. For instance, underground explosions produce seismic waves that can couple into the atmosphere in the form of infrasound. Likewise, atmospheric explosions produce infrasonic waves that can couple into the ground and travel as seismic waves [1]. AFTAC primarily relies on the seismic network for event detection, and, if infrasound station(s) near the detecting seismic stations also record arrivals, the combination of data between the two networks may refine location estimates. The Progressive Multi-Channel Correlation (PMCC) algorithm is a promising detection scheme for use on the infrasound network and is the subject of this research's evaluation.

PMCC is a correlation detector currently used by the international monitoring community to canvass recorded infrasound data for potential nuclear event detections. Specifically, the International Monitoring System (IMS) has tasked its data-processing arm, the International Data Centre (IDC), to use PMCC as part of its effort to ensure compliance with the Comprehensive Nuclear-Test-Ban Treaty (CTBT). As a correlation detector, PMCC assumes event-produced infrasound propagates through the atmosphere as a plane wave. Plane wave propagation implies that the wave front will reach each

1

sensor array's horizontally displaced sensors at predictable times given knowledge of the wave's angle of arrival and velocity. Of course, prior knowledge of both the angle of arrival, or azimuth, and propagating velocity is unknown. However, the time delay in plane wave arrival at one element relative to another can be calculated via cross-correlation of the two infrasound sensor elements' measured atmospheric pressure variations. The propagating signal's velocity and azimuth can be estimated from the computed arrival delays at all elements relative to arrival at a designated reference element. Generally, cross correlations are initially computed for the three possible sensor pairs of the three closest array elements.

In the ideal case, plane wave arrival means the three time delays for these pairs will sum to zero. In practice, how close the sum must be to zero in order to indicate plane wave arrival is set by the consistency, or PMCC's primary detection threshold. If additional elements can satisfy the consistency threshold when included with the initial sub-array of three, then the likelihood of detection is considered higher. On its own, consistency threshold satisfaction does not qualify as a detection. Rather, consistency satisfaction produces *elementary* detections in time-frequency space, and, if possible, elementary detections with similar signal attributes (i.e. angle of arrival, velocity, time of arrival, and frequency content) are grouped into a family. Large families generally signify higher likelihood of signal-of-interest (SOI) presence than do smaller families.

1.1 Problem Statement

The IDC does not use family size as a detection threshold. In fact, the organization only goes so far as to say that the *largest* and most stable families are

2

preserved for source localization [2]. Relatively little is known not only about the performance capability of PMCC, but also about how the choice of detection thresholds and PMCC-parameter settings affect this performance. In this regard, the consistency-dependent trade-off between the probability of detection (P_D) and the false alarm rate (FAR) is evaluated. This research additionally explores the feasibility of employing family size as the second of two detection thresholds – the first being the consistency. It is the job of this second threshold to eliminate false alarms and preserve true detections for further processing. Finally, globally-located infrasound stations experience a variety of wind noise conditions. Strong wind bursts introduce high-amplitude noise that can potentially render a station blind to SOIs. The limitations of PMCC detection capability are therefore investigated from a signal-to-noise ratio (SNR) perspective.

1.2 Research Contributions

This research develops a method to evaluate PMCC detection performance over a range of possible consistency thresholds. The initial step is to build a ground truth (GT) set of SOIs garnered from five infrasound stations – three along the Korean Demilitarized Zone (BRD, CHN, KSG), one located in Japan (I30), and one located in Russia (I45). Further, *optimum* family size detection thresholds are proposed for use by any agency using PMCC to monitor infrasound for nuclear events. Finally, methods to determine station-specific SNR failure levels are also described in detail. The following list summarizes this research's insights into PMCC detector performance:

- Receiver Operating Characteristic (ROC) curve analysis explores the consistency-dependent trade-off between P_D and FAR. Based upon this analysis, a recommended range of acceptable consistency thresholds is proffered.

- Both maximum a posteriori (MAP) and Bayes risk minimization approaches are used to determine optimum family size thresholds. These optimum thresholds depend on parameter settings such as the consistency threshold and the filter configuration used to filter data prior to running the PMCC detection algorithm. The proposed thresholds assist an analyst in making a decision as to whether or not a SOI is present based upon the size of the family in question.

- Synthetic signals are generated to determine signal-to-noise ratio (SNR) levels at which PMCC is blind to potential SOIs on various synthetic array configurations. These SNR detector limitations are then compared to the limitations of detection capability on fielded infrasound stations with geometries similar to those of the synthetically tested arrays. Detection "blindness" is defined as the SNR level at which $P_D \leq 10\%$.

1.3 Thesis Overview

Chapter II begins by discussing the background and current status of the CTBT. The nuclear test verification regime is introduced with particular emphasis placed upon the infrasound network and the PMCC detection scheme used to process the network's recorded data. Following this introduction is a comprehensive review of WinPMCC, the program that implements the PMCC algorithm [3]. Previous work characterizing the atmospheric propagation of infrasound, measures implemented to combat deteriorating

detection capability in the face of wind noise, and the effects of sensor array geometry on the ability to accurately estimate the azimuth and velocity of a propagating infrasound signal are also covered. Finally, infrasound signal detection methods other than PMCC are briefly described. Research objectives are included within the context of addressing some of the detection capability and performance challenges presented in Chapter II's literature review.

Chapter III reiterates these objectives and explains, in detail, the methodology tailored to achieve them. Chapter IV formally presents the results and analysis thereof. Chapter V summarizes the research contributions within the framework of recommending how monitoring agencies such as AFTAC or the IDC can streamline the use of PMCC in an operational setting and evaluate infrasound station performance. The document then concludes with recommendations for future work.

II. Literature Review

2.1 Background

2.1.1 Comprehensive Nuclear-Test-Ban Treaty (CTBT)

Over 2,000 nuclear tests were carried out between 1945 and 1996. 46 of those 51 years coincide with the Cold War, a period dominated by the Nuclear Arms Race and a series of conflicts with threatening nuclear undertones. Counted among these conflicts are the Berlin Blockade (1948-1949), the Korean War (1950-1953), the Suez Crisis (1956), the Cuban Missile Crisis (1962), the Vietnam War (1959-1975), the Yom Kippur War (1973), and the Soviet War in Afghanistan (1979-1989). Many attempts were made to organize a comprehensive nuclear test ban during the Cold War, but none came to fruition. Not until 1994, when memories of the not-so-distant past really catalyzed efforts to avoid future nuclear tensions, did negotiations successfully result in a treaty that would come into being two years later. The Comprehensive Nuclear-Test-Ban Treaty (CTBT) was first signed in Geneva in 1996. Presently, 183 countries have signed the treaty, and 157 countries have also ratified it [4].

In order for the CTBT to become law, all 44 countries holding nuclear technology must sign and ratify the treaty. 8 of these countries have yet to sign and/or ratify, including China, Egypt, India, Iran, Israel, North Korea, Pakistan, and the United States [4]. Without speculating as to the varied reasons why these countries have not taken their individual steps to help make the CTBT law, one possible reason why the United States has not ratified may have something to do with the fact that the measures put in place to verify whether a nuclear test occurs are not yet fully operational. Specifically, 337 global

facilities exist to monitor potential nuclear explosions. 15% of these monitoring facilities are not yet up and running [4].

2.1.2 International Monitoring System (IMS)

The Preparatory Commission for the Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO), headquartered in Vienna, tasked the International Monitoring System (IMS) with maintaining these facilities and their associated monitoring technologies. As the CTBT bans nuclear explosions by anyone, anywhere - on the Earth's surface, in the atmosphere, underwater, and underground - the IMS likewise uses 4 different sensor networks to ensure the detection of explosions by anyone, anywhere. Seismic, hydroacoustic, infrasound, and radionuclide stations comprise the IMS.

The seismic sensor network, consisting of 50 primary and 120 auxiliary stations, monitors shockwaves traveling through the Earth's crust and is therefore most conducive to detecting explosions underground. The 11 hydroacoustic stations monitor sound waves traveling through Earth's oceans, and the 60-station infrasound network monitors sound waves traveling through the atmosphere. Finally, the IMS's radionuclide component samples the atmosphere for radioactive particles with 80 globally-located stations [4]. Since shockwaves traveling underground and sound waves traveling through water and air are caused by a variety of events, the radionuclide component offers the only clear indication as to whether an explosion is nuclear in nature.

This research, however, will focus on the infrasound-monitoring component. The 60 infrasound stations are shown in Figure 1. Before delving into the infrasound network's operational details, what exactly is infrasound? What kinds of events produce

7

infrasound? To what does the term "infrasound station" refer, and what sensors are used to detect infrasound?

2.2 Infrasound

As its Latin prefix implies, *infrasound* describes a classification of sound waves with frequencies *below* the audible level of human hearing. Since the audible range refers to sound waves in the atmosphere with frequencies between 20 Hz and 20 kHz, infrasound signals exhibit frequencies less than 20 Hz. Infrasound typically travels at the speed of sound, or 343 m/s at 20℃. Many sources produce infrasound, including severe weather, bolides, ocean swell microbaroms, mountain associated waves, volcanic eruptions, auroras, earthquakes, rockets, and explosions [5] [6]. Examination of infrasound dating back to the World War II era illuminated the promise of garnering information not only about signal origin but also about the state of the atmosphere as a whole. Thereafter, infrasound detection and analysis was mainly developed to monitor nuclear explosions. When the CTBT was signed in 1996, infrasound research became more imperatively relevant. Shortly after the turn of the millennium, the IMS began construction on what would become the infrasound network shown in Figure 1 [2].

Each of the 60 stations on the map marks the location of an array of *microbarometer* sensors. These sensors measure atmospheric pressure and are most sensitive in the frequency range extending from hundredths of Hertz to a few tens of Hertz [7]. The frequency response, as illustrated in Figure 2, pertains specifically to the "MB2000" microbarometer, which is displayed in Figure 3. Individual microbarometers serve as sensor array elements.

Figure 1: The 60-station International Monitoring System (IMS) Infrasound Network as of 2010.
Figure copied from **[2].**

Infrasound arrays do not adhere to any worldwide standard of an "optimal" sensor configuration, partly because optimum sensor deployment is the subject of ongoing experimentation. However, the majority of the IMS infrasound monitoring stations have 7 or 8 array elements – microbarometers – with overall array apertures between 1.0 and 3.0 km [2]. Example configurations are shown in Figure 4. The reason sensors are arranged in such a way will be explored with the introduction of the detection method used to analyze the data recorded by these arrays. This research evaluates the capabilities, limitations, and implementation of this particular detection method.

Figure 2: MB2000 Frequency Response. *Figure copied from* **[7].**

Figure 3: MB2000 Microbarometer Sensor. *Figure copied from* **[7].**

Figure 4: Example 4-element centered triangle array (top), 8-element pentagon array with triangular sub-array (bottom left), and 9-element pentagon array with centered triangle sub-array (bottom right). *Figure copied from* **[2].**

The CTBTO requires that the infrasound network be mission capable at least 98% of the time. The requirement imposed on the infrasound network ensures it serves its role as an effective component of the IMS. "Mission capable" implies that at least 70% of the array elements at each station are correctly calibrated and transmitting their recorded data via satellite or virtual private network (VPN) to the International Data Centre (IDC), also located in Vienna [2]. The true litmus test for an effectively operating infrasound network is the ability to detect and locate any atmospheric nuclear explosion with a yield of at least 1 kiloton (kT) [8]. The academic community considers an explosive yield to

be detectable when there is a 90% probability of detection at two or more stations [9].

Detection ability hinges on the successful use of the Progressive Multi-Channel

Correlation (PMCC) algorithm, which is the detection scheme employed by the IDC to

process the IMS's station-level infrasound data that the IDC receives in real time.

2.3 The Progressive Multi-Channel Correlation (PMCC) Algorithm

PMCC begins by assuming that infrasound-producing events are far enough away

from surrounding sensor arrays that the arrays can treat the propagating infrasound

signals as plane waves. Infrasonic planar waves are represented using

$$p(\vec{r}, t) = e^{i(\omega t - \vec{k}\vec{r})} , \tag{1}$$

where $\vec{r}$ is the three dimensional position vector, $\vec{k} = \frac{2\pi f}{c}$ is the wave vector with

frequency f and phase velocity c, $\omega = 2\pi f$ is the angular frequency, and t is time [10].

Tabling the plane wave assumption for a moment, note that an infrasound signal in the

time domain $s(t)$ can be represented in the frequency domain by its Fourier transform

$$S(f) = A(f)e^{i\varphi(f)} , \tag{2}$$

where $A(f)$ represents the spectral amplitude, and $\varphi(f)$ represents the spectral phase.

Now, for a planar infrasound signal traversing a sensor array, the only difference between

the data recorded by any two sensor elements is a phase delay depending upon the

relative positions of the sensors $\Theta(\vec{r}_2 - \vec{r}_1)$ and the signal's incident azimuth and trace

velocity [11]. Of course, this ideal case assumes propagation free from attenuation and

background noise, and the following relations hold [12]:

$$A_2(f) = A_1(f) \tag{3}$$

and

$$\varphi_2(f) = \varphi_1(f) - \Theta(\vec{r}_2 - \vec{r}_1) \,. \tag{4}$$

Relaxing the two "ideal case" assumptions from which Eqns. 3 and 4 are born, Figure 5 illustrates what the three sensors of a triangular array might record as an infrasound signal passes. As opposed to the signal characteristics in Eqns. 3 and 4, background noise over time is characterized by rapid variations of both $A(f)$ and $\varphi(f)$.

Figure 5: Infrasound Arrival at a Small Triangular Array on Palmyra Island. The signal arrives at Sensor A at point 4343 and subsequently arrives at Sensors B and C respectively. The "time" units on the x-axis refer to the sampling rate (10 pts/s) [13].

13

Correlation is the basis for PMCC, and the cross correlation function determines the aforementioned signal arrival time delay between pairs of sensor elements $s_i(t)$ and $s_j(t)$. Cross correlations are performed within an analyzing time window of length W, where the channel data of sensor $s_i(t)$ is shifted over the channel data of sensor $s_j(t)$. The time shift at which the cross correlation is a maximum,

$$\Delta t_{ij} = \frac{1}{2\pi f}\left[\varphi_j(f) - \varphi_i(f)\right],\tag{5}$$

indicates the time difference of a signal's arrival between the two sensors. Equation 5 assumes the calculated delay is the same for all frequencies, i.e. dispersion is not a factor. This correlation operation is repeated for the two other sensor pairs in a three-sensor sub-array.

A plane wave produces a consistent set of time delays

$$\Delta t_{ij} + \Delta t_{jk} + \Delta t_{ki} = 0,\tag{6}$$

satisfying what is known as the *closure relation*. In the presence of background noise, the cross correlation operation may be less accurate due to random phase combinations, and the delays may not sum exactly to zero. The *consistency* of the set of time delays for n sensors of sub-array R_n is defined as the mean quadratic residual of the closure relation, expressed as follows:

$$c_n = \sqrt{\frac{6}{n(n-1)(n-2)}r_{ijk}^2}.\tag{7}$$

$r_{ijk} = \Delta t_{ij} + \Delta t_{jk} + \Delta t_{ki}$ and $i, j, k \in R_n$. If the calculated consistency is below an established threshold, a detection is declared on R_n [12].

Once a detection is declared, the time delays producing that detection are known and are subsequently inverted to obtain estimates for the propagating infrasound signal's velocity and azimuth [13]. Inversion is possible because a plane wave propagating from a fixed location at a specific velocity allows one to predict exactly when the signal will arrive at each sensor element as long as the array's relative position to the signal source is known. Conversely, knowledge of the time differences of arrival from sensor to sensor permits the trace velocity and azimuth to be estimated.

The element of PMCC yet to be discussed is the "P," or the algorithm's *progressivity*. Before broaching the subject, note that PMCC begins by determining the consistency on a set of delays for the *smallest* triangular sub-array. To reiterate, a detection is declared if the consistency is below an established threshold. The signal's trace velocity and azimuth, as determined by inverting the closure relation's time delays, is then used to "direct" the search for other sensors which may be added to the initial sub-array. Specifically, the value of the expected time delay for a pair of sensors, in which one of the sensors is outside the original consistency-evaluated sub-network, can be estimated. The computed time delay for this sensor pair corresponds to the correlation local maximum that is closest to the given estimate. As long as the detection criterion continues to be valid, i.e. the consistency threshold is met, the aperture of the network increases with each added sensor. As a result, velocity and azimuth estimates become more and more refined [11].

In addition to increasing signal attribute estimation precision, the progressive use of distant sensors helps reduce the number of false detections by removing correlated noise that may have been present on the original sub-network [2]. A potential pitfall,

however, stems from the ambiguity involved in the search for a correlation local maximum when adding distant sensors. Not only may two or more local maxima be close to the estimated expected time delay, but there may also be a local maximum closest to the estimate that has nothing to do with a coherent signal. For instance, noise on the far sensor of the sensor pair would conceivably produce numerous local maxima during a cross correlation operation. Since there is no minimum correlation value that must be met to qualify for consideration as a local maximum, this "seek and ye shall find" approach may unjustifiably add sensors to the original sub-array simply because there happens to be a local maximum in the area of the estimate. PMCC can therefore produce a high number of elementary detections that are false alarms. This high false alarm rate is tempered somewhat during the IDC's post-PMCC processing stage, the details of which will be discussed in Section 2.5.

The detection method just described, represented visually as the flow chart in Figure 6, is implemented at the station level and is therefore referred to as *station-level processing*. Station-level processing is the focus of the research described in Chapters III-V. *Network-level processing*, on the other hand, attempts to associate a signal-of-interest (SOI) detected on one array with an SOI on a neighboring array by comparing signal attributes garnered at the station level. These signal attributes, such as back azimuth, trace velocity, frequency content, and time of arrival, are used to determine whether SOIs recorded on multiple arrays were caused by the same event. If so, the IDC triangulates the position of this initial infrasound-producing event with each station's estimated back azimuth. These events are then added to standard event lists (SELs) and reported as part of an international bulletin [2].

16
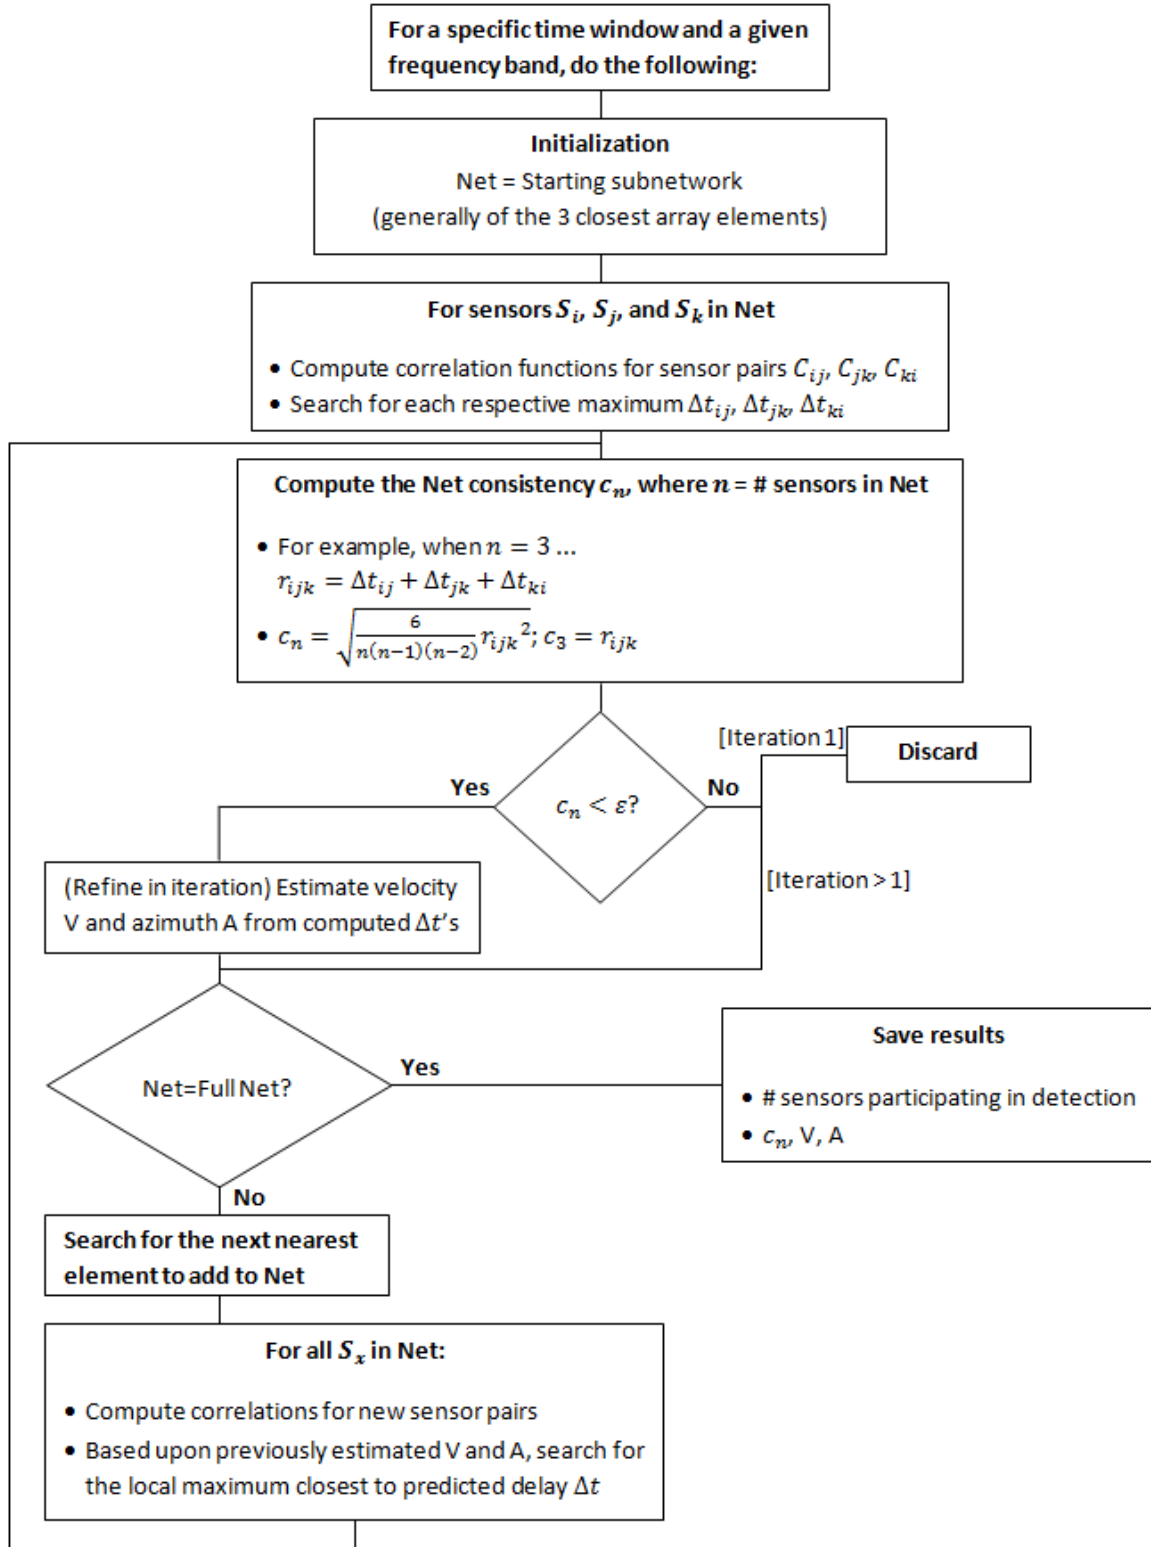
Figure 6: PMCC Flow Chart describing the steps necessary to produce an elementary detection. ε in the first "decision-making diamond" refers to the consistency *threshold*.

2.4 The Larger Picture and Implications

The introduction of nuclear treaty monitoring with infrasound via the PMCC algorithm has thus far disregarded the elephant in the room, namely satellite technology. Once satellite technology was established, atmospheric monitoring research efforts – especially US research efforts – went almost exclusively into furthering this technology. As a corollary, infrasound was largely neglected for 30 years. The reasoning was ostensibly simple. Why should the effort be made to process infrasound when satellites can "see" events in the atmosphere? The answer is twofold. First, the IMS does not operate satellites and needed another monitoring method. Second, the inclusion of infrasound in the IMS network permits expanded data fusion, or the synergistic combination of data from the seismic, hydroacoustic, infrasound, and/or radionuclide networks. In fact, seismic and infrasound stations are often collocated so infrasound can more easily be used as a discriminant for potential SOIs recorded by both infrasound and seismic arrays [14]. Data fusion therefore promises enhanced detection and localization capability and perhaps could even reveal events that may go undetected if only one method, including satellite technology, is used.

2.5 WinPMCC

The IDC's chosen algorithm with which to process the data it receives from the IMS infrasound network was introduced in the previous sections. WinPMCC is the implementation of this algorithm in programmatic form. Figures 7-9 and Figure 12 display the user options available when running the WinPMCC program, and these

18

options will be addressed as they relate to the operational use of the PMCC detection scheme.

A perusal of Figure 7 reveals some familiar terms, perhaps the most apparent being the "Threshold Consistency" under "Detection Parameters." Recall that this parameter serves as the threshold for detection, and the choice of its value presents an inherent trade-off between the probability of detection (P_D) and the probability of false alarm (P_{FA}), or the false alarm rate (FAR). The low probability of missing a true detection by setting a "high" consistency threshold will unfortunately be accompanied by

Figure 7: WinPMCC Parameter Settings Tab

a high FAR rate. Likewise, a "low" consistency threshold will reduce the FAR but at the expense of a higher probability of missing a true detection. *An objective of this research effort is therefore to investigate the consistency-dependent trade-off between P_D and FAR*, the details of which will be presented in Chapters III and IV.

Immediately adjacent to the "Threshold Consistency" in Figure 7 is "Threshold Nb of Sensors." This setting controls the minimum number of sensors that must participate in a detection. It ranges between three and the total number of sensors comprising the particular array [10]. A related parameter is "QLambda," which defines the maximum acceptable distance for integration of a far sensor into a sub-array, in accordance with the *progressive* aspect of PMCC. Q is a scalar multiplying the wavelength, λ, of a potential signal-of-interest (SOI). If the distance between sensors is not too large as compared to the wavelength, Q is generally set at a value, like 50, that will incorporate all sensors into the calculation [10]. Figure 8 shows an example sensor array and the corresponding network settings. The setting entitled "Sub Networks" indicates on which groups of sensors the closure relation will be initiated. These groups are generally smaller triangular sub-arrays to allow for the progressive use of distant sensors if the consistency threshold is met on any of the initial sub-arrays.

Returning to Figure 7, the "Filter Parameters" section presents the user with options for what filter configurations to use to filter the raw sensor waveform data. It also presents options for what analyzing time window length to use with each chosen filter. WinPMCC really begins station-level processing by filtering the data within the previously introduced sliding time window. Filtering increases the signal-to-noise ratio

Barycenter

	°long	min	sec	°lat	min	sec	z(km)
	126.547	0	0	38.2512	0	0	0

Sound Speed ☐3D 0.34 km/

5 sensors

Alias	°long	min	sec	°lat	min	sec	z(km)	#1	#2	#3	#4
BRD00_BDF	124.644	0	0	37.9657	0	0	0	☑	☑	☑	☑
BRD10_BDF	124.651	0	0	37.9709	0	0	0	☐	☑	☐	☑
BRD20_BDF	124.643	0	0	37.9611	0	0	0	☑	☐	☑	☐
BRD30_BDF	124.634	0	0	37.9683	0	0	0	☐	☐	☑	☑
BRD40_BDF	124.646	0	0	37.9625	0	0	0	☑	☑	☐	☐
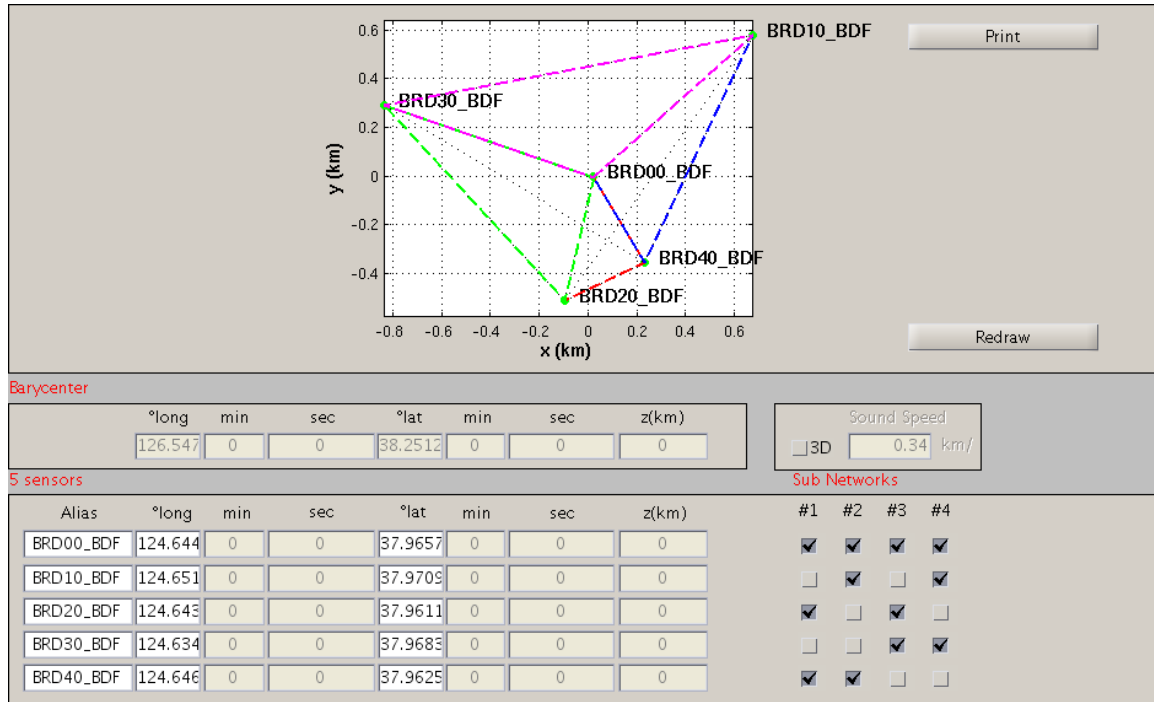
Sub Networks

Figure 8: WinPMCC Sensor Network Settings

(SNR) for SOIs whose frequency content lies within a specific filter's passband bandwidth. PMCC calculations then ensue post-filtering.

Multiple filters may be employed within a single time window, the idea being that additional filters may reveal signals with different frequency content that happen to simultaneously arrive at an array. "Nb of Bands" refers to the number of filters used, and overall filter configuration details are further specified in the "Window and Filter Parameters" dialog box in Figure 9. The top plot of Figure 9 vertically delineates the filters by bandwidths over which each filter's normalized magnitude response is unity. For example, "filter 20" has a normalized magnitude response that is unity between 4.55 Hz and 5.0 Hz, as can be visually confirmed in Figure 10's filter response plot.

Window lengths can also be individually adjusted for specific filters, with longer window

lengths generally applied to bandwidths containing lower frequencies. The filter

configuration recommended by Le Pichon and Cansi in the CTBTO's Technical PMCC

Documentation calls for 10 filters per decade spanning the following 3 bands:

- <u>0.5 Hz - 5.0 Hz</u>: This band is generally most appropriate for monitoring natural or

  man-made signals that propagate over distances of several hundred kilometers.

- <u>0.05 Hz - 0.5 Hz</u>: This band is generally most appropriate for monitoring remote large

  events, such as explosions or meteorite entries. Microbaroms are also often

  associated with this frequency range.

- <u>0.005 Hz - 0.05 Hz</u>: Detections in this band often point to atmospheric disturbances

  producing mountain associated waves and auroral infrasound [15].

To account for varying wavelengths, Le Pichon and Cansi further recommend using 30-

second window lengths for the highest of the above frequency bands, 90-second window

lengths for the middle frequency band, and 200-second window lengths for the lowest

frequency band [10].

As mentioned, PMCC calculations are consistently performed within these sliding

time windows and bandwidths. After canvassing the entire data segment for infrasound

arrivals, a list of elementary detections satisfying the consistency threshold remains, as in

Figure 11. These elementary detections are known as *pixels* within the WinPMCC

program, and, as is quite apparent, an almost constant stream of pixels is created in time-

frequency space. The seemingly innumerable detection list exists in no small part due to

Frequency (Hz)

4.55
4.1
3.65
3.2
2.75
2.3
1.85
1.4
0.95
0.5
0.455
0.41
0.365
0.32
0.275
0.23
0.185
0.14
0.095
0.05

30 40 50 60 70 80 90
Window length (s)

Band Index	WLength	Delta f	Order	Ripple
15	30	0.450000	2	0.01
14	30	0.450000	2	0.01
13	30	0.450000	2	0.01
12	30	0.450000	2	0.01
11	30	0.450000	2	0.01
10	90	0.045000	2	0.01
9	90	0.045000	2	0.01
8	90	0.045000	2	0.01
7	90	0.045000	2	0.01
6	90	0.045000	2	0.01
5	90	0.045000	2	0.01
4	90	0.045000	2	0.01
3	90	0.045000	2	0.01
2	90	0.045000	2	0.01
1	90	0.045000	2	0.01

Min :	30s	0.045Hz	2	0.01
Max :	90s	0.45Hz	2	0.01

Figure 9: WinPMCC Window Length and Filter Configuration Settings

23

Figure 10: Example of WinPMCC's Normalized Filter Magnitude Response

the ambiguity involved in the progressive search for distant sensors to add to initial sub-arrays, a PMCC pitfall more comprehensively addressed in Section 2.3.

WinPMCC's solution to this pitfall is to build pixel *families*, or group pixels that are similar in time-frequency-velocity-azimuth space and can therefore be associated with the same infrasound arrival [16]. In addition to eliminating pixels that cannot be associated with neighboring pixels, PMCC families help distinguish multiple arrivals that may exist in the same time window but in different frequency bands. Two pixels, P_1 and P_2, are grouped into a family if the weighted Euclidian distance between them is less than

$$d(P_1, P_2) = \sqrt{\frac{(t_2-t_1)^2}{\sigma_t^2} + \frac{(f_2-f_1)^2}{\sigma_f^2} + \frac{(V_2-V_1)^2}{\sigma_V^2 V_2 V_1} + \frac{(\theta_2-\theta_1)^2}{\sigma_\theta^2}}, \qquad (8)$$

where t_1 and t_2 are the times of arrival, f_1 and f_2 are the filters' center frequencies, V_1 and V_2 are the estimated trace velocities, and θ_1 and θ_2 are the estimated back-azimuths for P_1 and P_2 [10]. Whereas the azimuth indicates the angle of arrival *from* an infrasound source *to* a sensor array, the back-azimuth points *from* the array *to* the source.

24

Therefore, the true back-azimuth is 180° different from the true azimuth. Returning to

Eqn. 8, the σ's are weighting factors to allow for the comparison of quantities with

different units. The velocity weight σ_V is the only dimensionless parameter. This weight

is expressed in a percentage [12]. The weighting factors can be tuned independently;

default factors used throughout this research effort are shown in Figure 12.

Note also the parameters entitled "ThresholdFamMax" and "ThresholdFamMin."

"ThresholdFamMax" caps the maximum family size to a certain number of pixels to

obviate possible memory issues for infinitely growing families, as may be the case for

Figure 11: WinPMCC example of elementary detections (pixels) produced between 1300-1500 hrs on 25 August 2011 on the KSG Array, located in the Korean Demilitarized Zone (DMZ)

Figure 12: WinPMCC Families Settings used throughout this research effort. Relative to Eqn. 8, σ_θ = sigma_a, σ_V = sigma_v, σ_f = sigma_f, and σ_t = sigma_t.

microbarom detections, which can last between hours and days. Conversely, "ThresholdFamMin" specifies the minimum number of pixels that are necessary to constitute a family [10]. The post-family detection list is presented in Figure 13. Only the largest and most stable families are preserved for source localization in network-level processing [2]. As "large" is an ambiguous and relative adjective, *another primary objective of this research is therefore to determine the optimum family size* and quantify exactly what "large" should mean. Again, this research effort will be formally presented in Chapters III and IV.

After the building of families, the final element of station-level processing

26

Figure 13: Result of WinPMCC's example family-building procedure for the data recorded by the KSG Array on 25 August 2011 between 1300-1500 hrs

involves *detection categorization*, or the classification of PMCC families into either

"phase" or "noise" categories. Phases are infrasound detections that can be associated

with detections from other IMS stations, including other infrasound stations as well as the

seismic and hydroacoustic sensor network stations. The noise category is reserved for

coherent noise detections, or infrasound events which are of no concern to the IMS's goal

of CTBT compliance. Coherent noise may originate from a variety of sources, including

some infrasound-producing events already mentioned, such as large amplitude ocean waves (microbaroms), mountain associated waves, avalanches, and tornadoes. Overall, about 90% of infrasound detections are identified as noise with the current IDC algorithm [2]. Figure 14 demonstrates what a detection list might look like pre versus post-detection categorization, in which noise detections are removed.

Detections categorized as noise are labeled "N." The IDC categorizes detections other than noise with speeds greater than 2900 m/s as seismic in nature. All other detections are interpreted as infrasound arrivals and are named "I." Network-level processing combines all non-noise, station-level infrasound detections with detections from the seismic and hydroacoustic networks and attempts to localize events from these associations. Candidate events validated on two arrays are automatically reported in the SEL international bulletins [2]. This research does not investigate the network-level association process, but rather is concerned with evaluating and optimizing station-level processing.

Recall that the IMS strives to achieve a 90% probability of detection at two or more infrasound stations for explosions whose yields are at least 1 kT. Research by Green and Bowers cites external factors other than source yield that influence detection capability, including wind noise at the scale of a local array as well as seasonally shifting atmospheric wind directions. Approximately 80% of detected infrasound signals travel through the stratosphere [9]. Figure 15 compares the general shift in stratospheric wind directions from the Northern Hemisphere's summer to the Northern Hemisphere's winter.
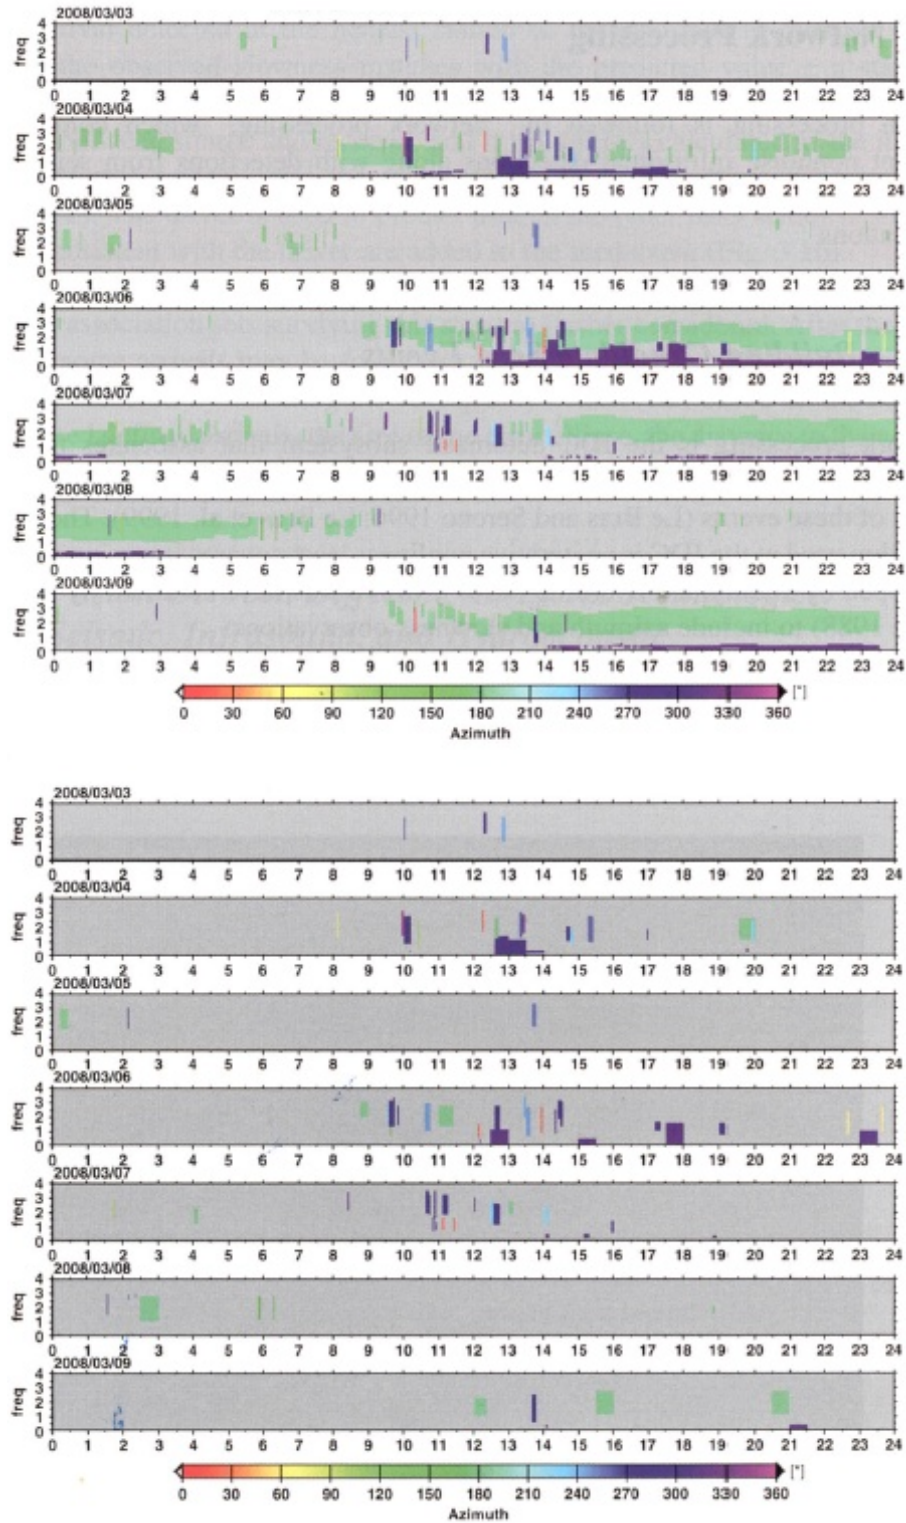
Figure 14: Example detection list before categorization (top) and after categorization and noise phase removal (bottom). *Figure copied from* **[2].**

2.6 Seasonal Winds and Atmospheric Propagation Considerations

When infrasound propagates in the same direction as stratospheric wind, i.e. downwind, the likelihood increases that infrasonic signals will refract into the troposphere (atmospheric level nearest Earth's surface), thereby increasing the likelihood of detection. On the contrary, when infrasound propagates upwind relative to the stratospheric waveguide, infrasonic signals are more likely to refract into the upper levels of the atmosphere, thereby decreasing the likelihood of detection [9]. Since stratospheric winds seasonally vary, the IMS infrasound network's detection capability varies seasonally as well.

Inclusion of stratospheric winds in detection capability models tends to lower the minimum yield that can still satisfy 90% probability of detection at two stations. The change in this 90% probability detection threshold with time is expounded upon in Figure 16. The caveat of Figure 16 is that its results are based upon the state of the infrasound network in October 2008, at which point only 39 of 60 total stations were

Figure 15: The dominant stratospheric wind direction is from east to west during the Northern Hemisphere's summer and from west to east during the Northern Hemisphere's winter. Infrasound from a given event is usually observed downwind. *Figure copied from* **[9].**
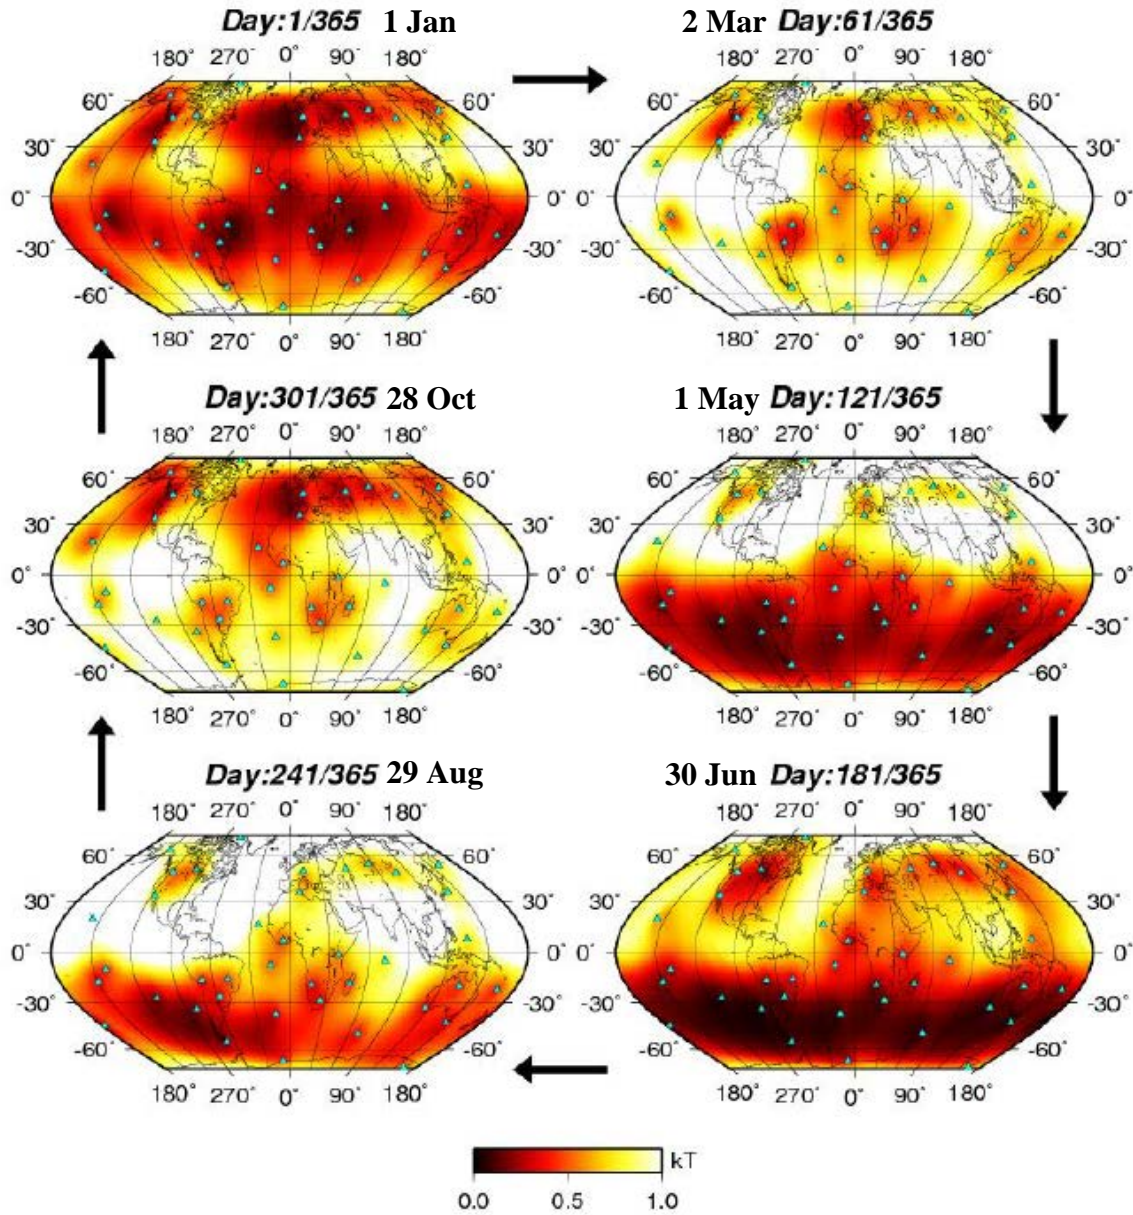
Figure 16: Stratospheric wind variability throughout the year and the related change in the 90% probability of detection at two stations for explosions with yields according to the color legend. These results are based upon the infrasound network in October 2008, when 39 of 60 stations were operational [9].

operational. For the completed 60-station infrasound network, models predict that 95%

geographical coverage at the 90% two-station detection probability level is achieved at

yields of ~0.6 kT during periods of high stratospheric winds and ~0.9 kT during periods

of low stratospheric winds [9]. In other words, detection capability models that account for seasonally-dependent stratospheric wind indicate that the infrasound network is more sensitive than what previous windless models had implied. For reference, the research expounded upon in Chapters III and IV involves infrasound-producing events that occurred in the Northern Hemisphere during the month of August.

Detection capability models can be further improved with a better understanding of the role wind direction plays in the relationship between explosive yield and recorded signal amplitudes. Researchers at the Los Alamos National Laboratory (LANL) did just that, establishing an empirical relationship between source yield and sensor-recorded pressure amplitude. The relationship,

$$P_{wca} = 5.95 \times 10^4 (SR)^{-1.4072} , \tag{9}$$

accounts for amplitude variability generated by stratospheric winds with climatological horizontal wind model HWM07. P_{wca} is the wind-corrected pressure amplitude, calculated from the peak-to-peak pressure of a stratospheric infrasound arrival, P_{raw}, using

$$P_{wca} = P_{raw} \times 10^{(-0.018)V_s} , \tag{10}$$

where V_s (m/s) is the component of the stratospheric wind velocity in the direction of propagation at an altitude of 50 km. SR from Eqn. 9 refers to the *scaled range* between the infrasound-producing source and the station recording the infrasound signal's arrival, defined as

$$SR = \frac{R}{\sqrt{2 \times Y}} . \tag{11}$$

R is the source-to-station range in kilometers, and *Y* is the explosive yield in kilotons [9]. The news surrounding advances in the understanding of atmospheric winds' impact on network performance is not exclusively optimistic, however.

Unfortunately, the increase in network sensitivity comes at the expense of diminished source localization ability. Since strong stratospheric winds reduce the likelihood of detection on arrays located upwind, often only arrays located downwind can participate in back-azimuth triangulation. Not only does the azimuthal separation of likely detecting stations decrease, but the distance to detecting stations will also likely increase [9]. Upwind stations that are potentially closer to the source than downwind stations may never record an infrasound arrival due to the increased probability that the signal refracts into upper atmospheric layers. These less than desirable stratospheric wind effects are more completely characterized in Figure 17.

2.7 Wind Noise and Deteriorating Detection Capability

In addition to atmospheric wind direction variability, the other primary external factor acknowledged by Green and Bowers as influencing detection performance is wind *noise* at the scale of a local array. Strong wind bursts introduce high-amplitude incoherent noise, potentially rendering an array blind to infrasound SOIs [2]. Therefore, *a third research objective is to determine detector limitations by investigating signal-to-noise ratios (SNRs) at which PMCC fails to register true detections*. Wind noise has long been known to hinder detection capability, which is why wind-reducing measures are built into the infrasound network wherever possible.

Figure 17: The percentage of Earth's surface across which (left) the azimuthal separation of the two most likely detecting stations is, at most, the angle indicated on the x-axis, and (right) the distance to the second most likely detecting station is, at most, the distance indicated on the x-axis. *Figure copied from* **[9]. Azimuthal coverage decreases with increasing stratospheric wind, and the distance between the two most likely detecting stations increases with increasing stratospheric wind. These results are based upon the infrasound network with 59 operational stations out of a possible 60.**

For instance, infrasound stations are often located in forests to minimize arrays' exposure to wind-generated background noise. Since forests do not ubiquitously inhabit the globe, other wind-reducing methods have been developed to diminish the "blinding effect" of wind-generated noise. Infrasound sensors are microbarometers sensitive to acoustic atmospheric pressure variations. Various pipe array designs, such as those shown in Figure 18, reduce wind noise by spatially averaging the micropressure field surrounding array elements. In addition to pipe arrays, designs for screened enclosures have also been introduced to further attenuate wind-generated noise. Design Version 5B, whose schematic can be seen in Figure 19, accomplishes this noise reduction while remaining virtually transparent to infrasonic signals in the monitoring passbands [2].

Figure 18: Examples of wind-noise-reducing systems employed throughout the IMS infrasound network at individual sensor elements. These pipe array designs reduce wind noise by spatially averaging the micropressure field surrounding a microbarometer sensor. The top rosette arrangements are most common. The bottom left design is less common, and the pipe array on the bottom right is designed to operate under snow cover at the Nuemayer Base in Antarctica. *Figure copied from* [2].

2.8 Sensor Geometry and PMCC Performance

Much consideration has been devoted to enhancing array detection opportunities, but optimizing the array configuration itself remains to be addressed. The importance of different array apertures and geometries becomes apparent when noting the variation in correlation coefficients over a range of possible azimuths.

Figure 19: Version 5B of the turbulence-reducing enclosure. *Figure copied from* **[2].**

Empirical observations on the reliability of infrasound detection reveal that certain array configurations exhibit azimuthally-dependent detection characteristics [2]. In this regard, Figure 20 compares three common sensor geometries.

Attempts to optimize array aperture, or sensor separation distances, must contend with the competing desire for accuracy between closure relations and source localization. Larger aperture arrays are more susceptible to cross correlation ambiguity than smaller aperture arrays, thereby leading to *less* reliable closure relations. The degree of signal correlation decreases as sensor separation increases due to the higher likelihood of path-altering effects [17]. For example, sensors separated by larger distances are more likely to be situated at different elevations than sensors located closer together. Therefore, the plane wave assumption may no longer apply, and the closure relation, which is based upon the horizontal distances between sensors, is less likely to be satisfied.
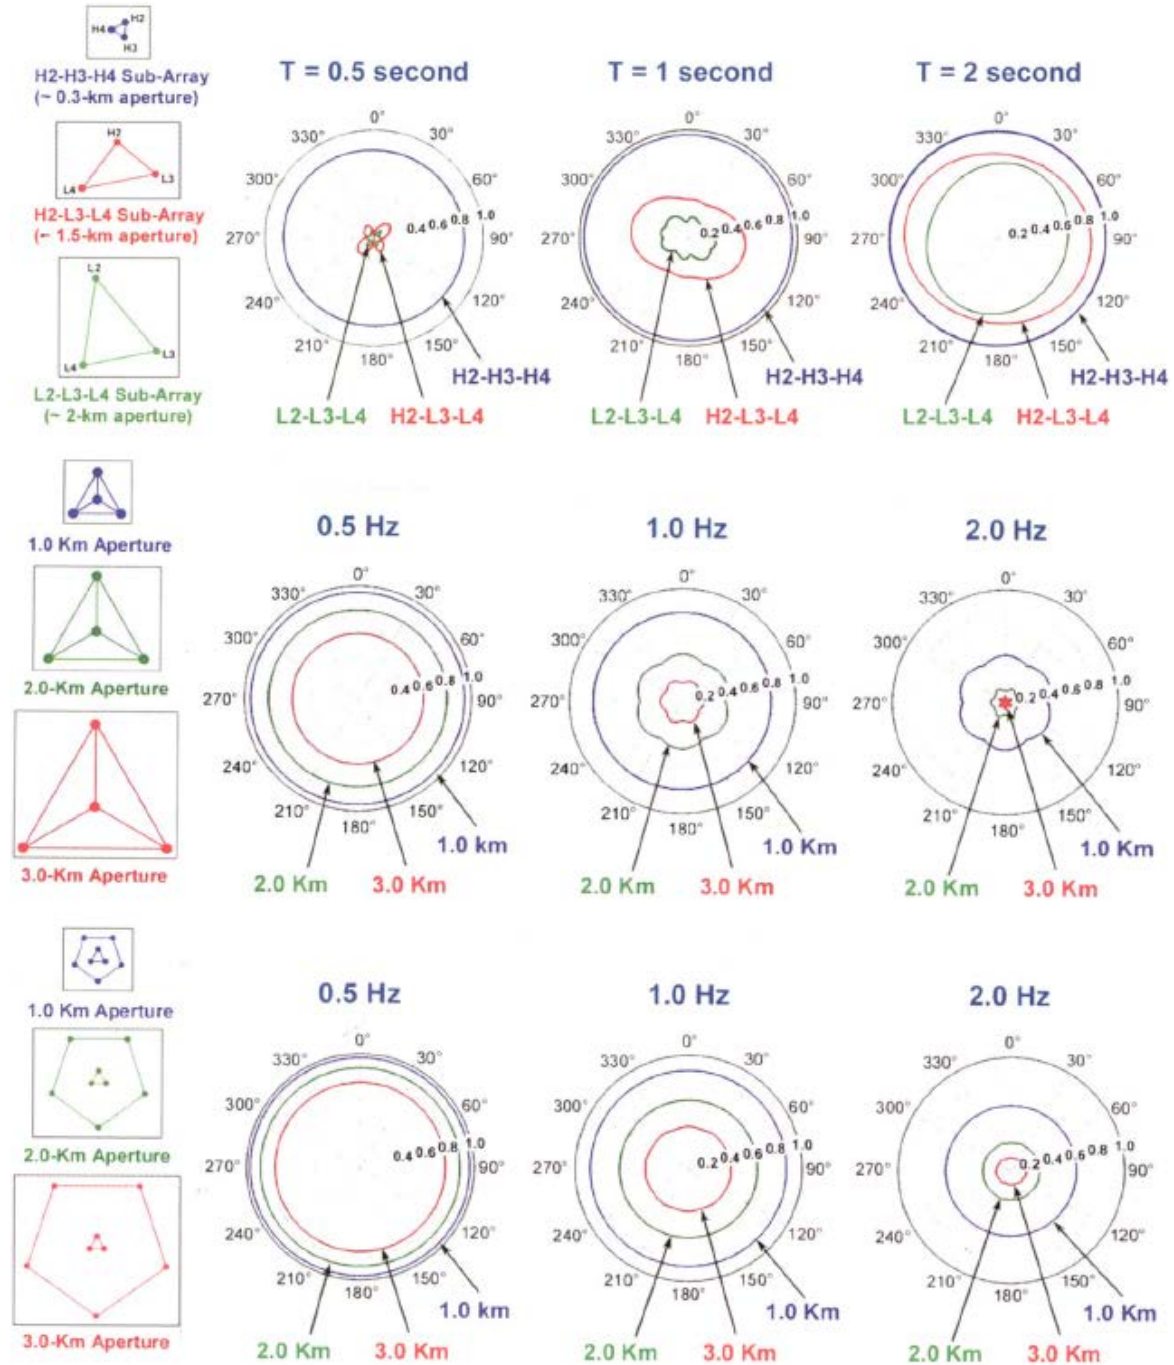
Figure 20: Predicted azimuthal variation of the array-averaged correlation coefficient for (top) triangle arrays, (middle) centered triangle arrays, and (bottom) pentagon arrays with triangular sub-arrays. The variation is based upon signal frequency and aperture size. *Figure copied from* **[2].**

In addition, larger aperture arrays imply larger time delays, exposing a propagating infrasound signal to more varying ambient wind conditions at each sensor element. Besides creating different noise environments, wind can alter signal propagation direction, further leaving the plane wave assumption for large aperture arrays on more tenuous ground. However, the progressive addition of distant sensor elements to consistency-evaluated sub-arrays leads to *more* reliable signal attribute estimation, specifically with regard to velocity and azimuth estimates. The benefits of increasing the array aperture were discussed when the progressive aspect of PMCC was introduced in Section 2.3, the most important benefit being improved source localization potential.

2.9 Signal Parameter Estimation: Azimuth and Trace Velocity

Estimation techniques, on the other hand, have not yet been covered here beyond the cursory claim that the time delays producing a detection can be inverted to obtain the propagating infrasound signal's velocity and back-azimuth. Szuberla and Olson propose incorporating the delay information into a matrix model and solving for estimates of trace velocity and back-azimuth with a least-squares approach [18]. The approach begins by noting the locations of the N sensor elements of an array in an (x_i, y_i), $i \in (1, N)$ fashion, where distances are arranged relative to an origin-defining sensor. The set of cross-correlation computed time lags, τ_i, indicates the plane wave's arrival at each sensor relative to a reference time. Finally, the unknown signal parameters velocity V and azimuth Θ are arranged in a two-element vector, and the matrix equation is presented as follows [19]:

$$\tau = Xf \, , \tag{12}$$

38

where

$$\tau = \begin{pmatrix} \tau_1 \\ \vdots \\ \tau_N \end{pmatrix}, \tag{13}$$

$$X = \begin{pmatrix} x_1 & y_1 \\ \vdots & \vdots \\ x_N & y_N \end{pmatrix}, \tag{14}$$

and

$$f = \begin{pmatrix} \frac{1}{V} \sin \Theta \\ \frac{1}{V} \cos \Theta \end{pmatrix}. \tag{15}$$

Above are N equations and two unknowns. If the computed time delays are not precisely accurate, Eqn. 12 is inconsistent. The least-squares method accounts for these time delay errors and solves Eqn. 12 in an approximate sense with

$$\hat{f} = (X^T X)^{-1} X^T \tau . \tag{16}$$

The terms in Eqn. 16 can be rearranged as

$$\hat{V} = \left(\hat{f}_1^2 + \hat{f}_2^2 \right)^{-1/2} \tag{17}$$

and

$$\widehat{\Theta} = \tan^{-1} \left(\hat{f}_1 / \hat{f}_2 \right) , \tag{18}$$

yielding least-squares parameter estimation equations for V and Θ [19].

2.10 Other Infrasound Signal Detection Methods

Not only had parameter estimation techniques not been previously covered, but there has also been no mention of any detection scheme other than PMCC. A whole host of alternative detection methods can theoretically process infrasound data, but the IMS prefers PMCC to these methods. One of the alternative methods involves using the sensor arrival time delays to align and overlap all of the sensor channel data into a single beam. Beam power is then plotted as a function of a two-dimensional wave number vector $\boldsymbol{\theta} = (\theta_1, \theta_2)$, which is nonlinearly related to the velocity V and azimuth Θ (f is signal frequency, and Θ is measured clockwise in radians) as follows [13]:

$$V = \frac{f}{\|\boldsymbol{\theta}\|} \tag{19}$$

and

$$\Theta = \tan^{-1}\left(\frac{\theta_1}{\theta_2}\right). \tag{20}$$

A maximum likelihood estimator operates on the plotted beam power, a graphical representation of which can be found in Figure 21 [20]. Based upon its use of time delays, this maximum likelihood approach, like PMCC, assumes the plane wave model holds. In fact, all of the detection methods discussed here will make this same assumption.

A second detection alternative also relies on beam-forming, but the detection statistic for this method is instead a function of the beam power divided by a noise power estimate. Division by the noise power estimate converts the beam into an F-statistic and, as a result, creates what is known as an F-detector [21]. The peak of the F-statistic in

Figure 21: Plotted as a function of $\theta_1, \theta_2 \in [-0.5, 0.5]$, the beam power is displayed as (left) a three dimensional surface plot and (right) a contour plot. The maximum likelihood detector determines that the beam power peaks at $\Theta = 225°$ and $V = 0.26\ km/s$, according to wave number Eqns. 19 and 20. *Figure copied from* **[13].**

Figure 22's example plots correspond to this method's best estimate for a signal arrival.

Other conventional methods, such as Capon's Method and Multiple-Signal Characteristic

(MUSIC) algorithms, share the same flaw as the beam-forming methods in that they

generally assume only one SOI is present at any one time. The IMS may have chosen

PMCC for its ability to discriminate simultaneously arriving signals whose frequency

Figure 22: Plotted as a function of $\theta_1, \theta_2 \in [-0.5, 0.5]$, the F-statistic is displayed as (left) a three dimensional surface plot and (right) a contour plot. The F-statistic peaks at $\Theta = 229°$ and $V = 0.24\ km/s$. Wave number Eqns. 19 and 20 apply just as in Figure 21. *Figure copied from* **[13].**

content lies in separate filter passbands. Despite this detection advantage, advances in the application of Fisher's F-statistic have led to at least one detection scheme that claims an advantage over PMCC.

2.11 InfraMonitor

Before delving into what this advantage is, it may be useful to clarify how the F-statistic may be used as a detection threshold. The variance of a data segment can be split into two separate variances, where both follow χ^2 distributions. One of these distributions is proportional to the total power in the sensor data, including SOI power and noise power, while the other is proportional to only the SOI power [22]. The F-statistic is based upon the ratio of these variances, where deviation of the ratio from unity indicates a SOI may be present [23]. The degree of deviation from unity allows statistically significant confidence levels to be assigned to detection declarations.

Arrowsmith et al. further improved the use of the F-statistic as the basis for a detection method by modifying its calculation to adapt to ambient noise conditions [24]. This improvement, incorporated into a program called InfraMonitor, precludes the requirement of applying post-detection categorizations. Recall that WinPMCC removes noise detections from continuous/repetitive sources in its post-processing phase categorization procedure. InfraMonitor, since it iteratively adapts to real ambient noise, should not flag infrasound produced by a continuous/repetitive source, such as microbarom ocean swells, as detections at all. An example of InfraMonitor at work can be found in Figure 23. The program recognizes that the correlated noise produced by a
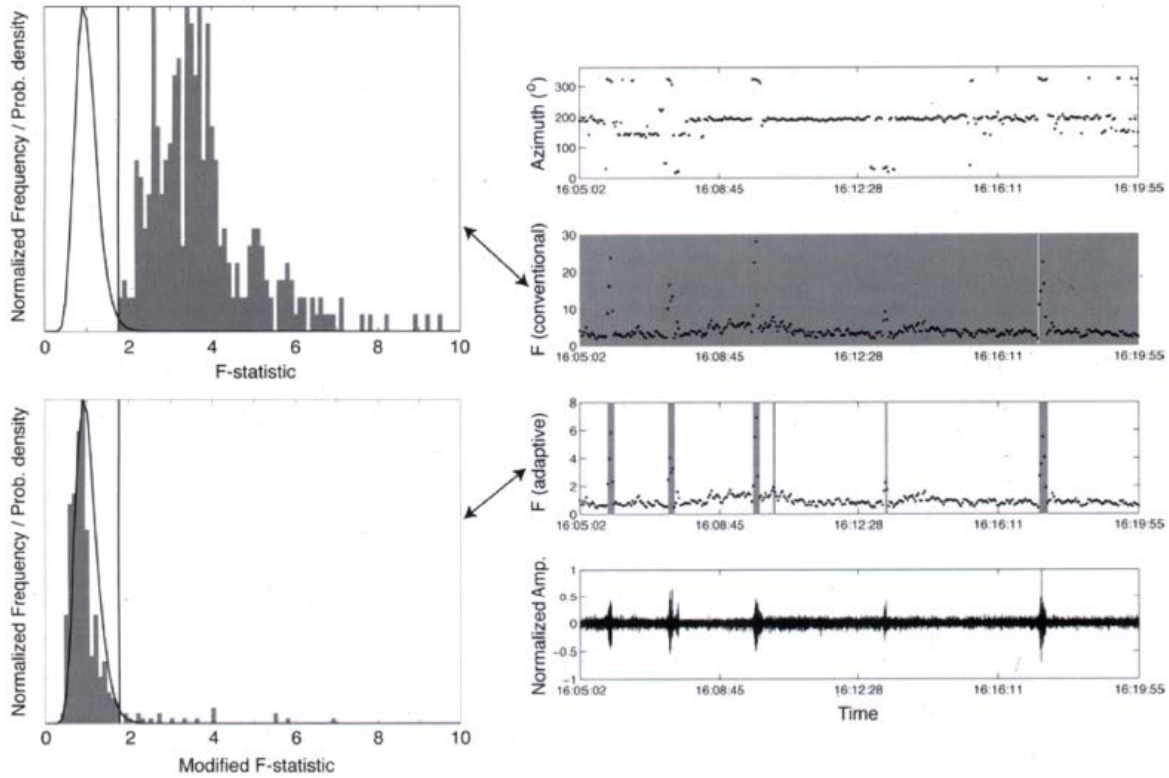
Figure 23: Example illustrating the difference between a (top) conventional F-detector and the (bottom) adaptive F-detector developed by Arrowsmith et al. The conventional F-detector flags nearly a constant detection (window marked in grey) for the correlated noise produced by a local wind farm. The adaptive F-detector recognizes infrasound produced by the wind farm as part of the ambient background, adjusts its detection threshold accordingly, and flags only other infrasound signal arrivals as detections (marked by grey vertical lines). *Figure copied from* [24].

local wind farm is just part of the ambient background and therefore ignores it as a source for detections.

2.12 Summary of Research Objectives

Ultimately, from the viewpoint of the research presented in later chapters, detection lists produced by InfraMonitor will be contrasted with those produced by WinPMCC. The collective results will help form the basis for a ground truth (GT)

detection set. Analysis of WinPMCC's performance, as applied to the GT set, will assist

in achieving the research objectives outlined throughout this chapter. Repeated here, the

primary objectives are to determine the consistency-dependent trade-off between P_D and

FAR, an optimum family size threshold(s), and the detection limitations of PMCC in

increasingly noisy environments.

III. Methodology

3.1 Chapter Overview

The previous chapter examined the development of the Progressive Multi-Channel Correlation (PMCC) algorithm and the details of its iterative detection scheme. Despite the availability of several other infrasound detection methods, some of which were introduced in Section 2.6, the International Data Centre (IDC) adopted PMCC and currently uses its algorithm to monitor infrasound-producing events. The International Monitoring System (IMS) keeps track of these events through IDC-submitted international bulletins called Standard Event Lists (SELs). These SELs assist the IMS in its mission to ensure compliance with the Comprehensive Nuclear-Test-Ban Treaty (CTBT). The Air Force Technical Applications Center's (AFTAC) mission is to use nuclear detection networks to detect nuclear tests carried out anywhere on the globe. Therefore, AFTAC needs to be cognizant of the performance capabilities and limitations of different detection and geolocation estimation algorithms used in the international community.

This research intends to further assist AFTAC, the IMS, and the IDC by offering a method by which these organizations can evaluate and ultimately improve infrasound station performance. The methodology presented shortly will explicate how PMCC consistency thresholds should be tuned and how family sizes can be optimized from a detection and estimation basis. Station geometry is also examined by determining how well PMCC performs in the face of deteriorating signal-to-noise ratio (SNR) conditions.

Before expounding upon these methods, a review of Figure 24 helps delineate where in the PMCC process the detection-discriminating "layers" fall. For instance, the

45
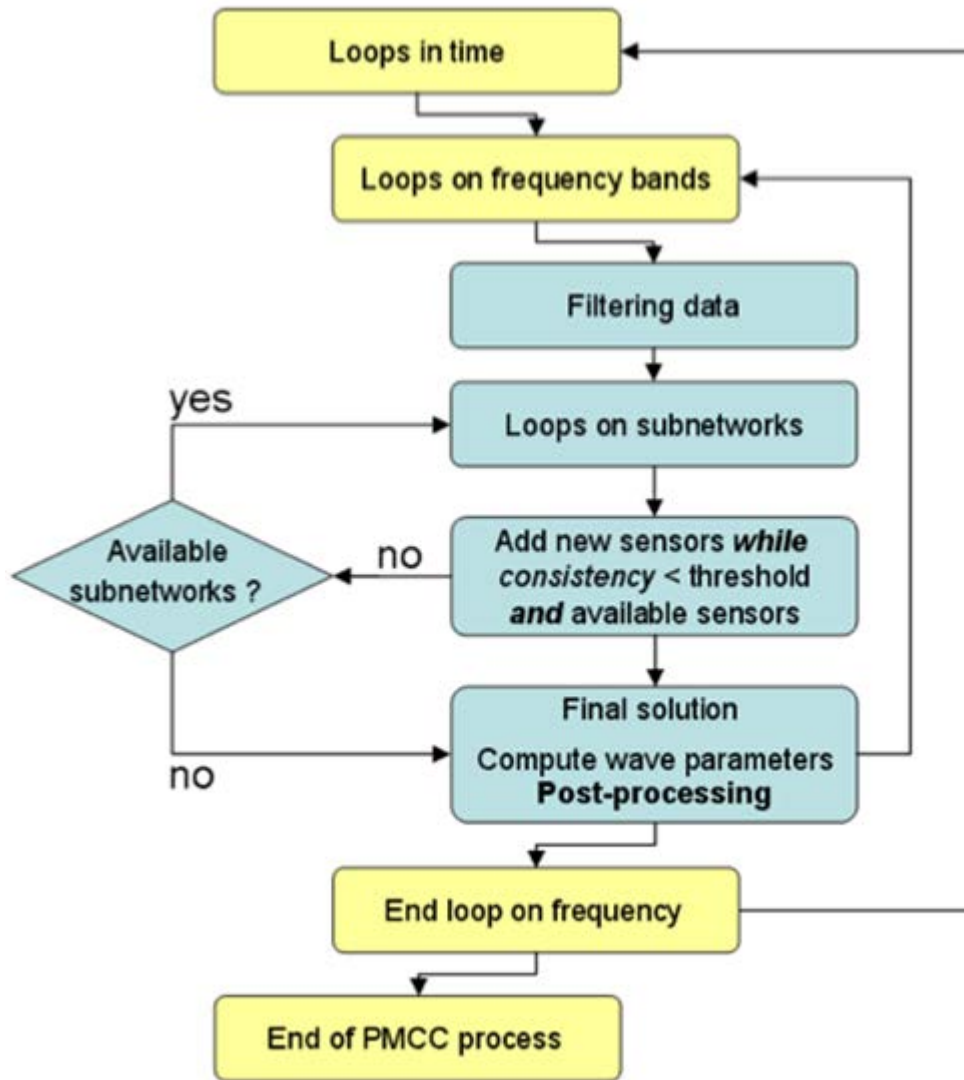
Figure 24: PMCC Flow Chart describing station-level processing. *Figure copied from* **[10].**

primary gatekeeper in differentiating coherent infrasound arrivals from incoherent noise

is the consistency threshold. Each successive detection-discriminating layer's job is to

remove false alarm detections from the previous layer's list while preserving the true

detections for subsequent data processing. In this regard, the layer following the list of

consistency-satisfied elementary detections is labeled "Post-processing" in Figure 24.

During the post-processing phase, elementary detections (pixels) with similar time-

frequency-velocity-azimuth signal attributes are grouped into families according to Eqn. 8. Not shown in Figure 24 are the final two detection-discriminating layers of PMCC, namely *phase categorization*, which identifies potential signals-of-interest (SOIs) from coherent noise, and *network-level processing*, which associates these SOI detections to detections produced by the same event on other infrasound/seismic/hydroacoustic stations. As mentioned, the consistency and family-building layers are the focus areas of this research.

3.2 Ground Truth Set of Detections

A prerequisite of this analysis is the establishment of a ground truth (GT) set of true detections. Three independent programs assist in building the set. Specifically, detections determined by WinPMCC are compared and contrasted with detections determined by Dr. Arrowsmith's InfraMonitor [24], the F-detector introduced in Section 2.11. A third program, SeaTools, proves useful in resolving whether detections flagged by either one or the other of these two detectors (but not both) are, in fact, *true* detections. SeaTools, not previously introduced, is a waveform analysis program initially developed by the Air Force Technical Applications Center (AFTAC) to review seismic data [25]. These three programs are used in concert to ensure the GT set is not biased towards any one program. Dr. Arrowsmith of Los Alamos National Laboratory (LANL) provided the time window of data analyzed, which consisted of detections during the month of August 2011 recorded by the 5 stations introduced in Section 1.2.

When used to canvass segments of sensor data, the WinPMCC program is run at a high (lenient) consistency threshold so as to not miss potential detections, even at the expense of a high false alarm rate (FAR), defined later in Eqn. 23. The burden in having to sift through a high number of false alarms to locate detections to add to the GT set is necessary to make sure the set includes infrasound SOIs that may have arrived under "less-than-desirable" conditions. One such condition could be high-amplitude incoherent noise at the scale of a local array. The fact that noise is present does not change the fact that a legitimate signal has arrived, but a lower consistency threshold may prevent WinPMCC from ever registering the signal's arrival as a detection. In other words, balancing the trade-off between the probability of detection (P_D) and the FAR is of little concern when the goal is to exhaustively include *all* true detections in the GT set. Table 1 specifies the settings used to run WinPMCC during this GT set-building process, including *family* settings and the chosen 10 second "high" consistency threshold (0.1 seconds is generally WinPMCC's default threshold).

At this point, it might be useful to clarify that a "detection" refers to a family, not merely a pixel. Likewise, InfraMonitor processes the same time segments of data as WinPMCC, and the two resulting lists are reviewed for common detections. Detections confirmed by both WinPMCC and InfraMonitor are then added to the GT set. An example of two detections confirmed in this manner is shown in Figure 25.

For those instances in which there is disagreement between WinPMCC and InfraMonitor as to whether a time window of data contains a detection(s), SeaTools's frequency-wavenumber (FK) analysis is called upon to resolve the dispute. The "FK Trend," as it is known in the AFTAC-developed program, not only keeps track of how

48

Table 1: WinPMCC settings used while building the GT set, including Filter Parameter settings, Detection Parameter settings, and Families settings. Example settings dialog boxes appear in Figure 7, Figure 9, and Figure 12. 90 second window lengths are used for the 10 filter passbands between 0.05 Hz and 0.5 Hz, and 30 second window lengths are used for the remaining 10 filter passbands between 0.5 Hz and 5.0 Hz. "Window Overlap" indicates the time shift for the sliding window lengths. "Ripple" refers to peak-to-peak passband ripple. "Threshold Nb of Sensors" refers to the minimum number of sensors that must participate in a detection. "QLambda" was explained in Section 2.5. "ThresholdFamMin" refers to the minimum number of pixels that must be grouped together before a family is created. "ThresholdFamMax" refers to the maximum family size. WinPMCC eliminates pixels whose estimated trace velocities are less than "VStoreMin" or greater than "VStoreMax."

Filter Parameters	
Nb of Bands	20
Freq Min	0.05 Hz
Freq Max	5.0 Hz
Window Overlap	50%
Order	2
Ripple	0.01 dB
Detection Parameters	
Threshold Consistency	10.0 sec
Threshold Nb of Sensors	3
QLambda	50
Families Settings	
ThresholdFamMin	5 pixels
ThresholdFamMax	300 pixels
VStoreMin	0.25 km/s
VStoreMax	0.45 km/s

the F-statistic varies within the time window analyzed, but also plots how the trace velocity and azimuth estimates vary as well [25]. When the computed F-statistic peaks concurrent with repeated velocity and azimuth estimates, as in Figure 26, a detection is confirmed. Table 2 specifies the settings used to run an FK Trend.

Random time windows of data from Figure 27's five stations are selected at different times of day and amount to a total of 45 hours. The final GT set, built from SOI arrivals within these time windows, contains 125 detections in the month of August 2011. To clarify, a "SOI" refers to a coherent infrasound arrival not produced by a continuous

Table 2: FK Settings. "Low Frequency" and "High Frequency" represent the range of expected SOI frequency content. Since slowness is the inverse of velocity, "Slowness Maximum" serves the same function as "VStoreMin" in Table 1. "Slowness Grid" specifies the number of points at which slowness is calculated. "Larger Window" is the overall window over which the FK Trend is computed. "Increment" indicates the time shift of the analyzing "Smaller Window" within the "Larger Window" and serves the same function as "Window Overlap" in Table 1.

Low Frequency	0 Hz
High Frequency	10.0 Hz
Slowness Maximum	5.0 s/km
Slowness Grid	81
Larger Window	120.0 sec
Smaller Window	10.0 sec
Increment	2.0 sec
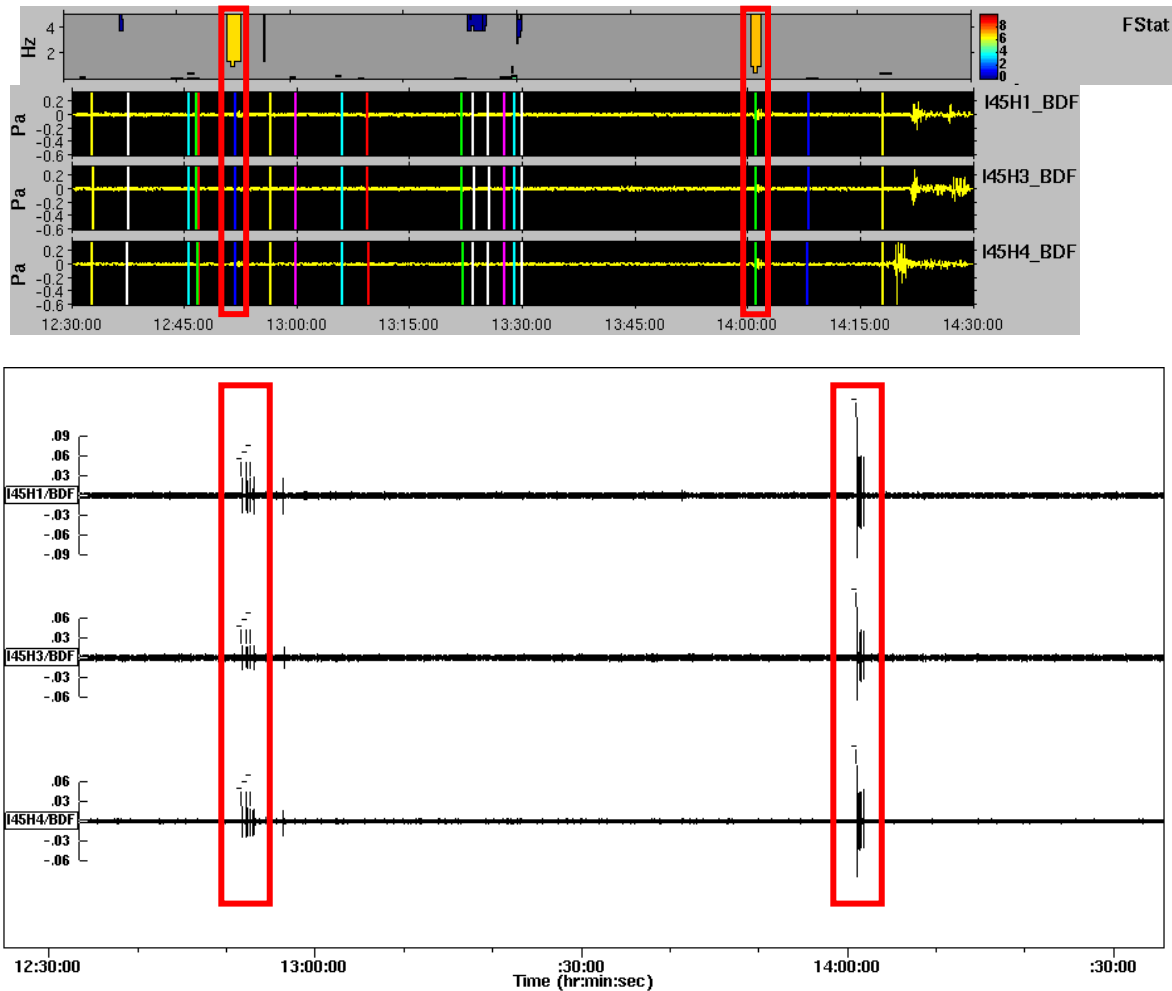
Figure 25: Two detections (in red) on the I45 Array on 25 August 2011 at 12:53 and 14:02. Confirmed by WinPMCC (top) and InfraMonitor (bottom)

50

or repetitive source. Recall that infrasound produced by such sources are labeled as "N" – for noise – and removed from subsequent analysis during WinPMCC's phase categorization process. Due to the arrays' proximity to the ocean, common repetitive sources are often ocean swell microbaroms.

3.3 Consistency Threshold and the Receiver Operating Characteristic

With a completed GT set, detector performance is now judged based upon how varying the consistency threshold affects the trade-off between P_D and FAR. Receiver

Figure 26: FK Trend, where an F-stat peak coupled with consistent azimuth and velocity readings confirms a detection on the CHN Array. This occurs at $t \approx 10{:}58{:}15$ for results presented here.

operating characteristic (ROC) curves generally depict this trade-off and are therefore commonly used to compare detector performance. Figure 28 illustrates three *example* ROC curves plotted on conventional axes. No actual data have been used to construct these curves. Rather, they are included for explanatory purposes in the case of a reader's unfamiliarity with ROC analysis. As the caption to Figure 28 explains, "steeper" ROC curves, or curves with greater area underneath them, imply increasing detector performance. Hypothetically, these *conventional* curves are created by plotting the fraction of true positive detections correctly classified (P_D) versus the fraction of true negative detections falsely classified (P_{FA}) at various binary decision-making thresholds. For a given threshold, P_D is calculated using [26]

$$P_D = \frac{\text{positives correctly classified}}{\text{total number of positives}}, \tag{21}$$

and P_{FA} is calculated using

$$P_{FA} = \frac{\text{negatives incorrectly classified}}{\text{total number of negatives}}. \tag{22}$$

In Eqn. 21, "positives correctly classified" refers to the number of true detections correctly identified (subset of a GT set) at a given threshold, and "total number of positives" refers to the size of the GT set. As mentioned, the size of the GT set for this research is 125. In Eqn. 22, "negatives incorrectly classified" refers to the number of instances in which detections are falsely declared (false alarms), and "total number of negatives" refers to the total number of instances in which detections should not be declared.
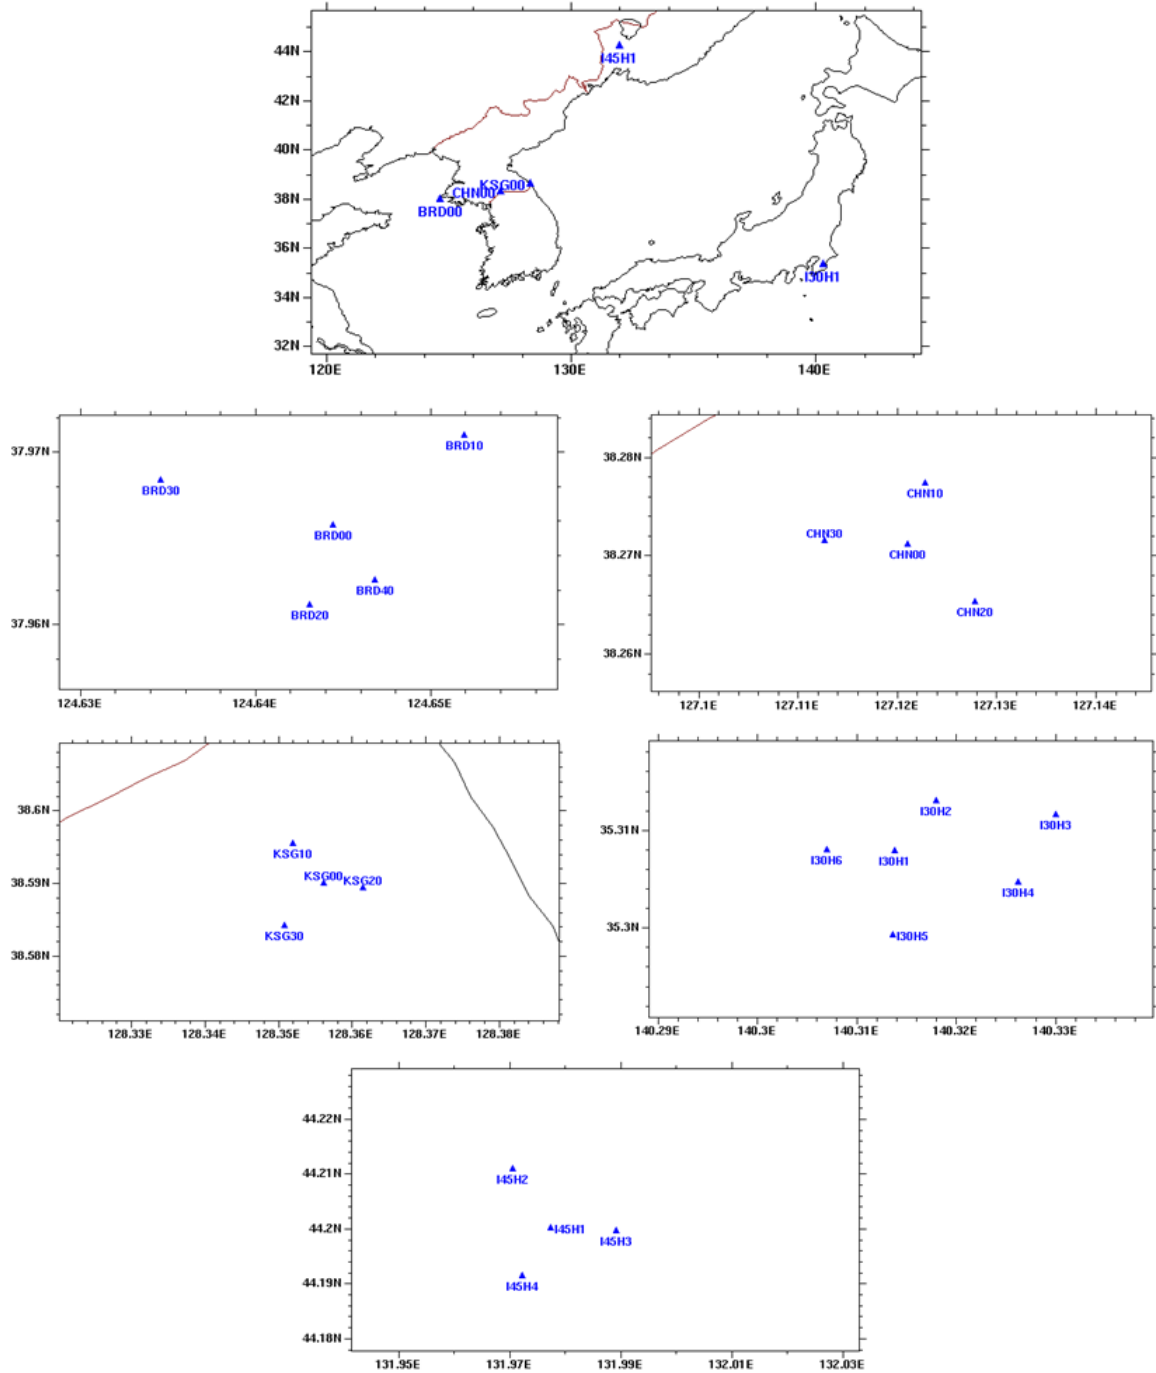
Figure 27: The GT set is constructed from SOI arrivals on the three arrays located along the Korean Demilitarized Zone (BRD, CHN, KSG), one array located in Japan (I30), and one array located in Russia (I45). The top chart shows the geographic locations of the stations, and the bottom five reveal the stations' array configurations.

Equation 22 conventionally applies to discrete binary tests, like a drug screening or pregnancy test. However, WinPMCC detector performance has been referred to in relation to a false alarm *rate*. For instance, note that the x-axis for the plot in Figure 28 is labeled as the *probability* of false alarm. This probability relies on the ability to assign a finite value to the denominator in Eqn. 22, the total number of negative detections. Considering that WinPMCC analyzes a *time* window of data within which the absence of detections cannot be quantified, WinPMCC's false alarm *rate* (FAR), computed as

$$\text{FAR} = \frac{\text{Total FAs}}{\text{Total Hours Analyzed}} \times 24 \left(\frac{hours}{day} \right), \tag{23}$$

Figure 28: Example conventional ROC curves. As the legend indicates, steeper ROC curves imply better detection performance. The linear ROC curve is labeled "worthless," because it is akin to random guessing. Any concave down curve, such as the blue curve, is "useful" because it represents a detector that outperforms random guessing. Any ROC curve that is concave up performs worse than random guessing.

is more appropriate. As Eqn. 23 suggests, this author chooses to quantify the rate on a

per day basis, but any length of time may theoretically be used. $\frac{\text{Total FAs}}{\text{Total Hours Analyzed}}$, for

instance, quantifies the FAR on a per hour basis. The graphical depiction of a detector's

P_D as plotted against its FAR is known as a *pseudo*-ROC curve, an example of which is

shown in Figure 29. As is the case for a conventional ROC curve, a "steeper" pseudo-

ROC curve signifies a better-performing detector.

Moving beyond this introduction to ROC analysis and into how it pertains to this

research, it must be reiterated that WinPMCC detections classified as "true positive"

Figure 29: Example Pseudo-ROC curves. As in Figure 28, the steeper pseudo-ROC curve implies better detection performance. However, unlike Figure 28, a pseudo-ROC curve that is concave up is not necessarily a "bad" detector. Rather, it is plagued by a FAR that initially increases more quickly than the probability of detection.

detections or false alarms do *not* refer to the elementary pixels solely satisfying the

consistency criterion. Rather, WinPMCC *families* produced as a result of the use of a

given consistency threshold that correctly identify a detection included in the GT set are

counted among true positive detections, while those that cannot be associated with GT set

detections are categorized as false alarms. The following consistency thresholds (sec) are

used to construct the pseudo-ROC curve: 1.0×10^{-7}, 1.0×10^{-6}, 0.01, 0.1, 0.5, 1.0,

and 10. WinPMCC's default threshold, 0.1 sec, is the test group's median. Note that the

most lenient threshold, 10 sec, corresponds to the threshold used to establish the GT set.

In discussing the building of the GT set, it was mentioned that the use of such a lenient

threshold implied an increased burden in having to sift through a high number of false

alarms. The false discovery rate (FDR)

$$\text{FDR} = \frac{\text{number of false alarms}}{(\text{number of false alarms+number of true positive detections})} \tag{24}$$

provides some insight into the burden on an analyst whose job is to review the list of

WinPMCC detections [26]. Therefore, the FDR is determined for each of the thresholds

used to construct the pseudo-ROC curve. With the exception of the varying consistency

threshold, the WinPMCC settings used throughout this pseudo-ROC-building process are

the same settings listed in Table 1.

3.4 Optimum Family Size

Table 1 and Figure 12 reveal that one of the parameters that can be adjusted prior

to running the WinPMCC program is the minimum number of pixels that must be

grouped together before a family is created. Recall from Section 2.5 that only the largest

and most stable families are preserved for source localization in network-level processing [2]. However, no clarification is proffered as to what constitutes a "large" or "stable" family. Therefore, this author proposes a method to determine the optimum family size and quantify exactly what "large" should mean.

The optimum family size is determined with a *maximum a posteriori* (MAP) approach in which the goal is to *minimize* the total number of false alarm and missed detection categorization decisions. Specifically, the solution to this approach indicates how many pixels must comprise a family before it is more likely than not that the family represents a true infrasound SOI arrival.

The first step in this approach requires organizing the GT set according to the number of pixels comprising each detection. The frequency with which a particular family size appears as a detection is then recorded. This process is repeated until every one of the 125 detections in the GT set are accounted for, and a probability histogram is created to visualize the distribution of family sizes. The histogram is then curve-fit with the probability density function (pdf) that best characterizes its distribution, as in Figure 30. This pdf is known henceforth as the conditional distribution of family sizes given the detection is a true event family, or $p(z|T)$, where z is the number of pixels in the family. The conditional distribution of family sizes given that the detection is a non-event family, $p(z|R)$, is determined in much the same way as was $p(z|T)$. Over all time periods from which the GT set was built, families that are neither in the GT set nor removed as coherent noise from a repetitive source are considered to be members of H_R, the rejection (or null) hypothesis. These noise detections are organized in the probability histogram in Figure 31, and overlaid onto this histogram is the exponential pdf best

57

Figure 30: Probability Histogram of GT set family sizes. Overlaid on the histogram is the lognormal pdf, $p(z|T)$, that best fits the data.

characterizing its family size distribution.

Note that this categorization marks a departure from the categorization used during the creation of the pseudo-ROC curve. Families not in the GT set had been considered false alarms. Those same families (post-repetitive source removal) are now considered SOI rejections. Why the difference? Well, it depends on the perspective of the decision-making entity. On one hand, the WinPMCC program producing a family is its way of declaring a detection. That declaration is either a true detection or a false alarm, and the pseudo-ROC curve is built based upon the accuracy of these declarations

58

Figure 31: Probability Histogram of non-event family sizes. Overlaid onto the histogram is the exponential pdf, $p(z|R)$, that best fits the data.

at various consistency thresholds. On the other hand, from the perspective of an analyst reviewing the list of families produced by WinPMCC, a decision has not yet been made as to whether a SOI is present or not. This optimality discussion exists to assist the analyst in making a decision based upon the *size* of the family in question.

The decision criteria are arranged in a likelihood ratio test (LRT)

$$\frac{p(z|T)}{p(z|R)} \underset{H_R}{\overset{H_T}{\gtrless}} \frac{P(R)}{P(T)} = \gamma \, , \tag{25}$$

which serves an integral part in determining the MAP family size threshold. The LRT

forms an inequality between the ratio of the true detection and rejection conditional

probability density functions and the ratio of the a priori probabilities of the presence

$P(T)$ or absence $P(R)$ of a SOI, notated conventionally as γ. $P(T)$ is the fraction of the

total number of families that are true detections, and $P(R)$ is the fraction of the total

number of families that are rejections. $P(T) + P(R) = 1$. H_T and H_R are the two

possible hypotheses [27]. Based upon the number of pixels composing the family in

question, choosing H_T implies that the presence of a SOI is more likely, and choosing H_R

implies SOI absence is more likely.

Returning to the optimality discussion, the terms of Eqn. 25 are now rearranged to

position the true event and non-event likelihood functions on either side of the inequality,

as follows:

$$\boldsymbol{P(T) \cdot p(z|T)} \underset{\boldsymbol{H_R}}{\overset{\boldsymbol{H_T}}{\gtrless}} \boldsymbol{P(R) \cdot p(z|R)} \,. \tag{26}$$

The true event likelihood function, $P(T) \cdot p(z|T)$, is simply the true detection conditional

probability density function scaled by the a priori probability that any single family is a

member of the GT set. Likewise, the non-event likelihood function, $P(R) \cdot p(z|R)$, is the

rejection conditional probability density function scaled by the a priori probability that

any randomly chosen family is a SOI rejection. The graphical intersection of these

likelihood functions marks the MAP threshold family size z_{t_MAP}. This intersection is

shown later in Figure 50 in Section 4.3.1. Families with more pixels than z_{t_MAP} are

more likely to be true detections, and families with fewer pixels than z_{t_MAP} are more

likely to be rejections. SOI presence and absence categorization decisions based upon

this threshold minimize the probability of error P_{error}, defined as $P_{error} = P_{FA} + P_{MD}$, where P_{FA} refers – as it did before – to the probability of false alarm, and P_{MD} refers to the probability of missed detection.

3.5 SNR Stress Tests and Detector Failure

Another basis of comparison for which to assess the performance of different stations and their various geometries is to "stress test" array configurations under deteriorating SNR conditions. A "failure SNR level," defined as the SNR at which $P_{MD} \geq 90\%$ ($P_D \leq 10\%$), are determined using both synthetic and real data. Recall from Section 2.5 that PMCC calculations on a time window of data do not commence until after that waveform data are filtered according to the filter configuration established in WinPMCC's "Window and Filter Parameter" settings. Therefore, *failure SNRs* are synonymous with *post-filtered SNRs*.

3.5.1 WinPMCC Filter Duplication

The only way to ascertain the post-filtered SNR is to duplicate WinPMCC's data filtering operation. Since WinPMCC often employs multiple filters, as Figure 9's 10-filters-per-decade configuration demonstrates, the question arises as to which filter to duplicate. The answer is the one that maximizes the post-filtered SNR, for such a filter gives WinPMCC the best chance of detecting a SOI, should a SOI be present.

For the synthetic data SNR stress tests, the synthetic SOI is the Pierce Blast shown in Figure 32. The filter most appropriate to duplicate for the purposes of determining the post-filtered SNR depends upon the Pierce Blast's power spectral density (PSD), which is plotted in Figure 33. The PSD reveals that the Pierce Blast's signal

61

Figure 32: Pierce Blast – the synthetic SOI used for the synthetic SNR stress tests. The vertical red line denotes the last point at which the amplitude is above 0.001 Pa. SOI power is computed between 0 sec and this vertical red line.

Figure 33: Pierce Blast Power Spectral Density

Figure 34: WinPMCC-produced magnitude response for "filter 16," whose lower and upper cutoff frequencies are 2.75 Hz and 3.20 Hz respectively. These cutoff frequencies mark the range over which the filter's passband is unity (or 0 on a dB scale).

power peaks at 3 Hz. The appropriate filter to duplicate is therefore "filter 16," whose lower and upper cutoff frequencies are 2.75 Hz and 3.20 Hz respectively. These cutoff frequencies mark the range for which the filter's magnitude response is unity (or 0 on a dB scale), as shown in Figure 34.

Having reviewed the procedure for selecting which filter to duplicate, all that remains is how to duplicate the filtering operation itself. To ensure the filter is recreated exactly, the duplicate filter must have the same transfer function, i.e. the same filter coefficients. The cutoffs and coefficients for each of WinPMCC's 20 filters are listed in the filter initialization file, the location of which is specified in the "File Settings" tab, as exemplified in Figure 35. The initialization file appears in Figure 36. The "ForwardCoeffs" within the file refer to the coefficients in the numerator of the transfer function, and the "ReverseCoeffs" refer to those in the denominator. Note that $a(1)$, the denominator's first filter coefficient, in

63

$$H(z) = \frac{b(1)+b(2)z^{-1}+\ldots+b(2n+1)z^{-n}}{1+a(2)z^{-1}+\ldots+a(2n+1)z^{-n}} \tag{27}$$

is always 1.00. As a result, the initialization file lists the "ReverseCoeffs" beginning

with $a(2)$. $H(z)$ is the filter's transfer function, and b and a are the "Forward" and

"Reverse" coefficient row vectors specifying the transfer function's zeroes and poles,

respectively, in descending powers of z. Equation 27 is based upon a MATLAB

convention, where n is the filter's order dictated by the "WinPMCC and Filter

Parameter" settings in Figure 9. Moreover, knowledge of the fact that WinPMCC's

source code is MATLAB permitted a trial and error process which revealed that

WinPMCC's filter configuration consists of chebyshev filters using MATLAB's *cheby1*

and *filtfilt* commands. Successful filter duplication is confirmed in Figure 37's

comparison of the WinPMCC and MATLAB-replicated filter magnitude and phase

responses.

Files

PMCC File

/home/dean/winpmcc/winpmcc/pmcc.ini   Select...

Referenced Files

Filter File

/home/dean/winpmcc/winpmcc/filters_I45_05-5Hz_6Nov.ini   Select...

Result File

/home/dean/winpmcc/winpmcc/results_I45_05-5Hz_6Nov.txt   Select...

Families File

/home/dean/winpmcc/winpmcc/families_I45_05-5Hz_6Nov.txt   Select...

Bulletin File

/home/dean/winpmcc/winpmcc/bulletin_I45_05-5Hz_6Nov.txt   Select...
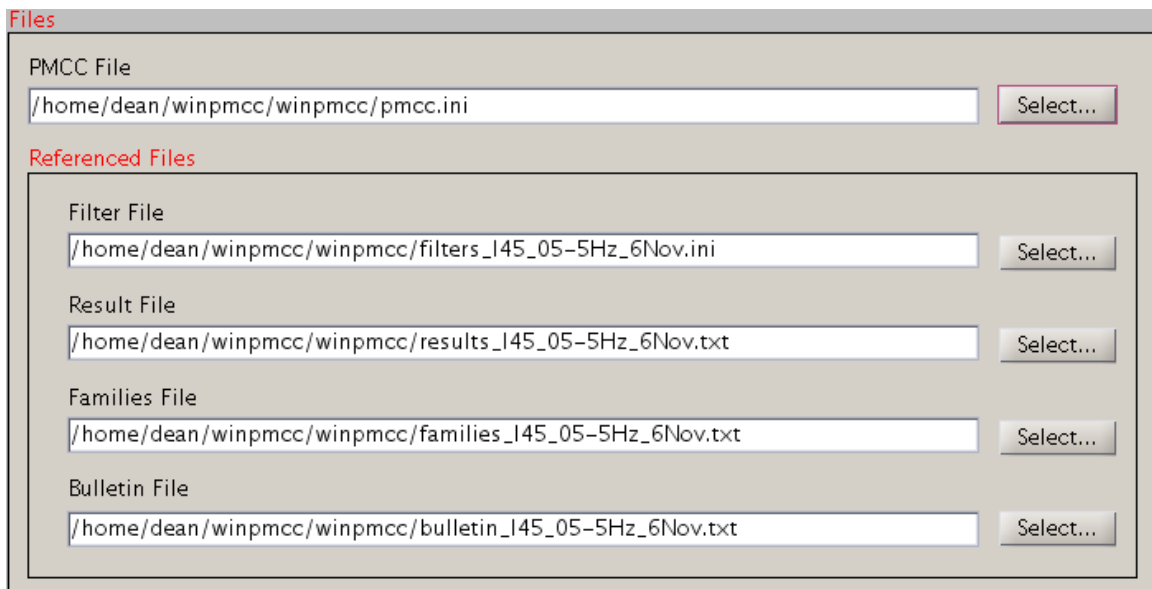
Figure 35: WinPMCC "Files Settings." "Filter File" refers to the filter initialization file.

3.5.2 Synthetic Data SNR Stress Test

Having confirmed the ability to duplicate WinPMCC's filtering operation, the synthetic data SNR stress test further requires the creation of a synthetic waveform. The synthetic SOI used in this analysis is the Pierce Blast in Figure 32, and additive white Gaussian noise (AWGN) is modeled with MATLAB's *randn* function. The SOI power and noise power are computed in the time domain using the numerical approximation for average signal power

$$P_{Av}^T \approx \sigma^2 + \mu^2 \,, \tag{28}$$

where σ^2 refers to the variance of the signal, and μ^2 refers to the squared mean of the signal [28]. The word "signal" is used here interchangeably to refer both to the SOI and the noise signal. P_{Av}^T, as stated, is the average signal power in the time domain. Since $Var(X(t)) = E[X(t)^2] - E[X(t)]^2$, Eqn. 28's simplified equivalent is

$$P_{Av}^T \approx E[X(t)^2] \,, \tag{29}$$

where $E[\cdot]$ is the expected value operator.

To ascertain the time-averaged SOI power P_{SOI}, the time range over which Eqn. 29 is applied is delineated between 0 sec and the vertical red line in Figure 32. This line denotes the last instance in which the SOI's amplitude exceeds 0.001 Pa. Equation 29 is also applied to realizations of *randn* noise vectors to determine the time-averaged noise power P_{noise}. Since *randn* generates pseudorandom numbers from the standard normal distribution, i.e. zero-mean with unity variance, the terms in Eqn. 28 reveal that P_{noise} approximately equals one. SNR manipulation now begins with the

65

```
Fmin = 0.001250, 0.002375, 0.0035, 0.004625, 0.00575, 0.006875, 0.008, 0.009125, 0.01025,
0.011375, 0.0125, 0.02375, 0.035, 0.04625, 0.0575, 0.06875, 0.08, 0.09125, 0.1025, 0.11375

Fmax = 0.002375, 0.0035, 0.004625, 0.00575, 0.006875, 0.008, 0.009125, 0.01025, 0.011375, 0.0125,
0.02375, 0.035, 0.04625, 0.0575, 0.06875, 0.08, 0.09125, 0.1025, 0.11375, 0.125

ForwardCoeffs =
0.00012804989516256121639091458686, 0, -0.00025609979032512243278182917372, 0,
0.00012804989516256121639091458686
0.00012804989404074105033282116 0396, 0, -0.0002560997880814821006 65642320791, 0,
0.00012804989404074105033282116 0396

                                  ⋮

0.01111219497543944192269327331246, 0, -0.02224389950878883845386 54662492, 0,
0.01111219497543944192269327331246
0.01111219497543937687056292418 22, 0, -0.02224389950878753741125 84836439, 0,
0.01111219497543937687056292418 22
0.01111219497543942127948390918 846, 0, -0.02224389950878842558967 81837691, 0,
0.01111219497543942127948390918 846
0.01111219497543941364670061489 051, 0, -0.02224389950878827293401 22978101, 0,
0.01111219497543941364670061489 051
0.01111219497543942873879485588 873, 0, -0.02224389950878857477589 71177745, 0,
0.01111219497543942873879485588 873

ReverseCoeffs =
-3.96825248605065716844819689868, 5.90550644977896688203600 20618, -
3.90625178088103819362686408567, 0.96899783067409495096645 58847772
-3.96783389847631218216861270776, 5.90467590720178847618626 605254, -
3.90583973339578394146087703 12, 0.96899783067409595015281 010413

                                  ⋮

-3.29311905501378898719622156932, 4.41336738738252876146361813881, -
2.81221143601544065759867407905, 0.73083740142174358211946 1279945
-3.16753664561021430756682093488, 4.20934084181018519643657782581, -
2.70496833851813089921733990195, 0.73083740142174558052090 5605226
-3.02613426768994075999330561899, 3.98910595915400367772463141591, -
2.58421552708796209429920054 39, 0.73083740142174513643169 5755164
-2.86961814233804712870323783136, 3.75705701677760205364904 9736314, -
2.45055609046198030753771490708, 0.73083740142174391518636 8667491
-2.69876997485188541148204421916, 3.51782401426341451156076800 544, -
2.30465757832183815168036744581, 0.73083740142174435927557 8517554
```

Figure 36: Filter initialization file for filter configuration specified in Table 1. The cutoff frequencies listed under "Fmin" and "Fmax" are normalized to the sampling frequency, which is 40 samples/sec. For example, the 16th "Fmin" frequency is 0.06875, and the 16th "Fmax" frequency is 0.08 (see highlights). These values correspond to the lower and upper cutoff frequencies for "filter 16." Unnormalized, 0.06875 refers to 2.75 Hz $\{(0.06875) \times (40)\}$ and 0.08 refers to 3.20 Hz $\{(0.08) \times (40)\}$. Also highlighted, the "ForwardCoeffs" refer to the coefficients in the numerator of the transfer function of the filter, and the "ReverseCoeffs" refer to those in the denominator. ReverseCoeffs are listed beginning with the 2nd pole.

multiplication of a noise vector's amplitude at each time sample by a scale factor (SF)

$$SF = \sqrt{\frac{P_{SOI}}{(P_{noise})(SNR_{desired})}}, \qquad (30)$$
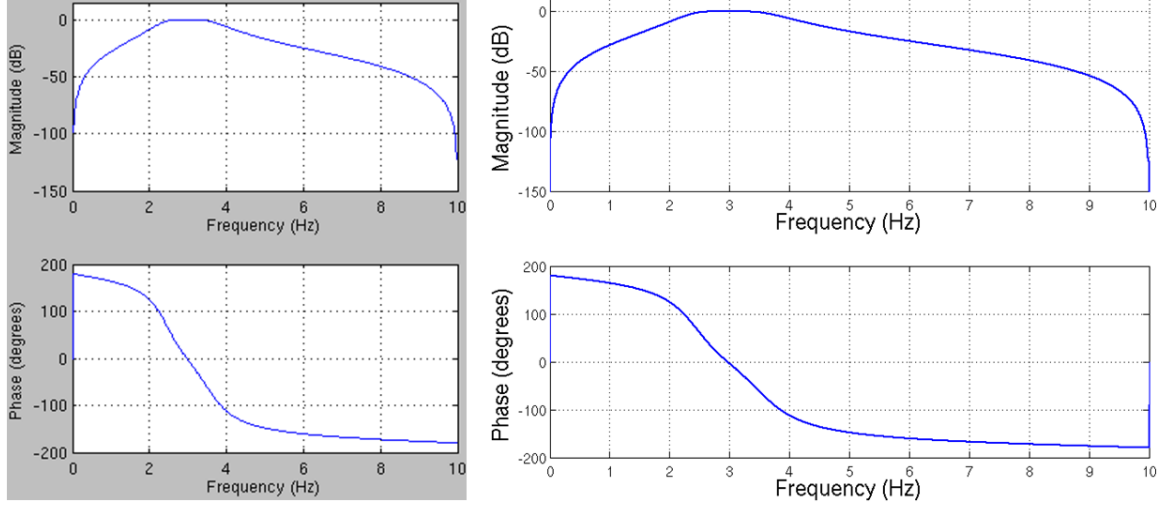
Figure 37: Proof of successful filter duplication. The magnitude and phase responses for "filter 16" produced by WinPMCC (left) and duplicated in MATLAB (right).

where $SNR_{desired}$ is the desired *pre*-filtered SNR expressed on a linear scale. As this author prefers to initially express desired SNRs in decibel (dB) units, conversion from $SNR\ dB_{desired}$ to the $SNR_{desired}$ appearing in Eqn. 30 is achieved with

$$SNR_{desired} = 10^{SNR\ dB_{desired}/10} \ . \tag{31}$$

Figure 38 shows 5 different noise vector realizations, each multiplied by the SF necessary to achieve the stated $SNR\ dB_{desired}$. The SOI is zero-padded so as to match the length of the noise vector to which it is added. The result of the addition

$$Waveform_{Synthetic} = SOI + (SF)(Noise) \tag{32}$$

creates a synthetic waveform at a desired pre-filtered SNR. An example synthetic waveform is illustrated in Figure 39, where the pre-filtered SNR, as expressed on a logarithmic scale, is 0 dB.

67

The synthetic waveform's post-filtered SNR is determined by separately filtering the SOI and noise with the MATLAB-duplicated filter prior to the addition in Eqn. 32. The post-filtered SOI power $P_{SOI,post-filter}$ and noise signal power $P_{noise,post-filter}$ are also separately computed with Eqn. 29, and the post-filtered SNR, on a dB scale, is

$$SNR\ dB_{post-filter} = 10 \times \log_{10}\left(\frac{P_{SOI,post-filter}}{P_{noise,post-filter}}\right). \qquad (33)$$

Before the *randn* AWGN signal is added, the post-filtered SOI is overlaid onto the pre-filtered SOI in Figure 41, demonstrating little change after applying "filter 16." Figure 42 presents this same pre-filtered versus post-filtered comparison for the noise signal. The post-filtered waveform, the one upon which the PMCC algorithm operates, is shown in Figure 40.

Of course, PMCC calculations occur on an array of sensor data and not merely a single sensor's data. The synthetic array in this analysis is the pentagon array with centered triangle sub-array appearing in Figure 43. The location of the Pierce Blast within each sensor element's data time window is shifted to reflect a plane wave arriving from a certain azimuth relative to the array and traveling at a certain velocity. For a 40° azimuth and 300 m/s trace velocity, the Pierce Blast arrives on each of the array's nine elements as shown in Figure 44, in which WinPMCC's successful detection is also shown. Noise is added in SNR decibel level decrements until WinPMCC fails to detect the SOI at least 90% of the time. To determine if the deterioration of detection ability accelerates for configurations with fewer sensor elements, the synthetic SNR stress test is repeated on five synthetic arrays whose configurations are manipulated so as to resemble

Figure 38: Example *randn* noise signals which, when added to the Pierce Blast SOI, will produce synthetic waveforms with pre-filtered SNRs of -10 dB, -5 dB, 0 dB, 5 dB, and 10 dB respectively.

those from which the GT set was built. These five synthetic arrays and their actual array

counterparts are portrayed in Figure 45.

3.5.3 Real Data SNR Stress Test

The real data SNR stress test proceeds in a similar manner as the synthetic stress

test. Only now, a detection from the GT set is the SOI, and instead of modeling noise, a

time window within which there are no produced families *is* the noise. Of course, the

69

Figure 39: Example pre-filtered synthetic waveform at an SNR of 0 dB

Figure 40: Post-filtered version of the waveform in Figure 39. The post-filtered SNR = 12 dB .

70

Figure 41: Synthetic SOI (Pierce Blast) pre-filtering (blue) and post-filtering (red)

Figure 42: Synthetic noise pre-filtering (blue) and post-filtering (red) with "filter 16." When added to the Pierce Blast SOI, the pre-filtered SNR = 0 dB, and the post-filtered SNR = 12 dB.

71

Figure 43: Synthetic 9-element pentagon array with centered triangle sub-array

time window containing the detection has its own noise, but it is impossible to separate

this noise from the SOI. It is therefore also impossible to manipulate the amplitude of

that noise without altering the SOI. An infrasound SOI arrival on the BRD station is

depicted in Figure 46. In order to manipulate the SNR, a detectionless window of data,

composed entirely of noise, is overlaid onto the time window containing the detection for

each of BRD's five sensor elements.

To illustrate this process, note first that the time window containing the detection

in Figure 46 is $t = [03\colon 14\colon 15\ 03\colon 14\colon 30]$. A 15 second detectionless window from

elsewhere in the sensor data, say $t = [03\colon 12\colon 00\ 03\colon 12\colon 15]$, is then added to the

Figure 44: Example WinPMCC detection of Pierce Blast on the synthetic array in Figure 43 at a post-filtered SNR of 12 dB. The 9-element sensor data, as shown, have been filtered by a bandpass butterworth filter with -3 dB cutoffs of 2.0 Hz and 4.0 Hz. However, the filter configuration specified by Table 1 performed the filtering during the WinPMCC detection process.

15 second detection window. The amplitude of the added noise can now be increased until WinPMCC misses the detection. Since the noise profile of each element may differ, as is clear from the sensor data in Figure 47, so does the element-specific SNR prior to manipulation. Figure 47 consists of an 8-minute (480 sec) excerpt from the BRD array, where the detection from Figure 46 is located at the 6-minute mark ($t = 360$ sec). The SOI is much more visually apparent in Figure 48, the post-filtered version of the waveforms.

73
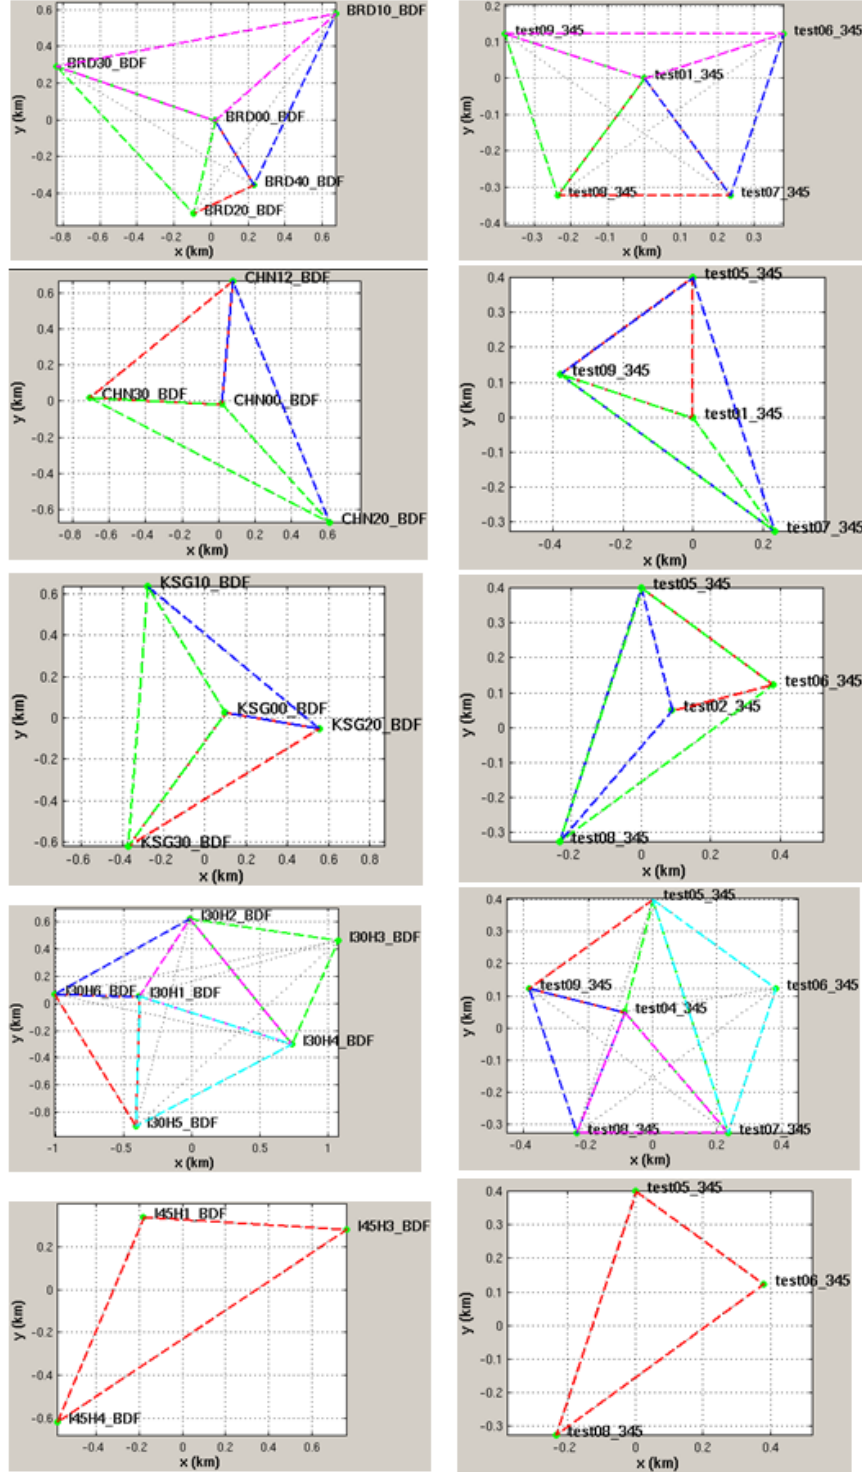
Figure 45: The real BRD, CHN, KSG, I30, and I45 arrays (left) are pictured opposite their synthetic counterparts (right). SNR stress tests are conducted on these synthetic arrays to determine if the deterioration of detection ability accelerates for configurations with fewer sensor elements than the 9-element array in Figure 43.

74

Figure 46: WinPMCC detection of a real infrasound SOI on the BRD array on 25 August 2011 at 03:14:15. The sensor data, as shown, have been filtered by a bandpass butterworth filter with -3 dB cutoffs of 2.0 Hz and 4.0 Hz. However, the filter configuration specified by Table 1 performed the filtering during the WinPMCC detection process.

Just as the ability to view and manipulate the post-filtered waveform required the use of a MATLAB-replicated filter in the synthetic stress test, so too is it required for the real data stress test. The choice of which filter from the WinPMCC configuration in Figure 9 to duplicate, however, no longer requires a PSD computation. Instead, the WinPMCC "families file," whose location is specified in Figure 35, provides information about the frequency content for each pixel comprising a family. The distribution of SOI power in the frequency domain is thus revealed, and the decision as to which filter to duplicate depends on which filter's passband captures the majority of SOI power. In other words, which filter accounts for the greatest number of produced pixels?

Figure 47: Eight-minute excerpt of pre-filtered data recorded by the BRD array. The detection from Figure 46 is located at the six-minute mark ($t = 360\ sec$).

Coincidentally, the answer to this question is "filter 16," the same filter used for the synthetic SNR stress tests.

To establish the same post-filtered SNR on each array element, the multiplicative scale factor by which to multiply the pre-filtered overlaid noise varies due to the unique noise profiles at each sensor. Since it is impossible to separately compute the SOI power and noise power in the 15-second detection window, an SNR proxy is used. This proxy, the short-term-power-average (STA) over long-term-power-average (LTA), STA/LTA, compares the post-filtered signal power in the 15-second detection window with the post-filtered signal power in the detectionless portion of the larger window. Just as in the synthetic stress test, "signal" is used here interchangeably to refer both to the SOI and the

76

Figure 48: BRD sensor data of Figure 47 post-filtering with "filter 16"

noise. The STA approximates the SOI power and is calculated with Eqn. 29. The LTA, computed with the same equation, approximates the noise power. The post-filtered SNR, expressed in dB, is calculated with Eqn. 33, where STA replaces $P_{SOI,post-filter}$, and LTA replaces $P_{noise,post-filter}$ as follows:

$$SNR\ dB_{post-filter} = 10 \times \log_{10}\left(\frac{STA}{LTA}\right). \tag{34}$$

To achieve a station-wide post-filtered SNR for which to test WinPMCC's detection ability, the appropriate multiplicative scale factors by which to multiply the overlaid noise vary, as mentioned, for each sensor element. For the BRD SOI arrival depicted in Figure 46 through Figure 48, the five multiplicative factors are determined

77

through an iterative trial-and-error process until the desired post-filtered SNR is achieved for each of BRD's five sensor elements. Noise is added in SNR decibel level decrements until WinPMCC fails to detect the SOI at least 90% of the time.

3.6 Summary of Evaluation Approaches

Synthetic and real data SNR stress tests can be performed by monitoring agencies, such as the IDC, to evaluate arrays' PMCC detector limitations in the face of increasing ambient noise. Use of synthetic arrays permits any number of geometries to be evaluated using synthetic data in the manner outlined by this chapter. The results of the synthetic SNR stress tests for these geometries can be compared against the results obtained on real arrays with similar geometries. Array performance can further be judged with ROC curve analysis by viewing the consistency-dependent trade-off between P_D and FAR. Finally, optimum family size thresholds can be determined on an array-by-array basis, thereby assisting an analyst with the decision for whether an individual detection should be considered for further processing.

IV. Analysis and Results

4.1 Chapter Overview

The previous chapter developed methods to determine optimum detection thresholds for the Progressive Multi-Channel Correlation (PMCC) algorithm used by the International Data Centre (IDC) to perform infrasound station-level event detection. Statistical detection theory via a maximum a posteriori (MAP) approach points to optimum family size thresholds of grouped detection pixels before detections should be considered for network-level processing. An additional approach is developed utilizing Bayes cost criteria. Optimum family sizes for these approaches are determined based upon the consistency threshold and filter configuration employed by the WinPMCC program. The consistency threshold is further explored insofar as it presents a trade-off between the probability of detection (P_D) and the false alarm rate (FAR). Additionally, synthetic signals at various signal-to-noise ratios (SNRs) are generated to determine SNR failure levels for the PMCC algorithm on certain synthetic array configurations. Detector limitations for these synthetic signals/arrays are compared to the SNR detector limitations of fielded infrasound stations with similar configurations.

4.2 Consistency Threshold and the Receiver Operating Characteristic

As potential detections initially require consistency threshold (c_n) satisfaction, this chapter first explores the consistency-dependent trade-off between P_D and FAR with Figure 49's Receiver Operating Characteristic (ROC) curve. Recall that c_n is defined by Eqn. 7 and further elaborated upon in the context of its purpose within the PMCC algorithm in Figure 6. These results suggest that threshold consistencies below

Figure 49: Pseudo-ROC curve presenting the trade-off between the probability of detection P_D and the false alarm rate (FAR) for all five GT stations in Figure 27 at the following consistency thresholds (sec): 1.0×10^{-7}, 1.0×10^{-6}, 0.01, 0.1, 0.5, 1.0, and 10.

1.0×10^{-6} sec cause WinPMCC to miss an unacceptable number of true detections (greater than 20%), while thresholds above 1.0 sec increase the FAR without an appreciable increase in P_D. In building the ground truth (GT) set, described in Section 3.2, the high FAR expected to result from using a 10-second threshold was ignored in favor of exhaustively including all potential SOI arrivals.

The use of such a threshold may elicit concern due to the increased possibility of the inclusion of false alarms in the GT set. However, recall that only detections confirmed by at least two of the three detection schemes (WinPMCC, InfraMonitor, and FK Trend) were added to the GT set. Also, as Figure 49 reveals, all true detections that satisfied the 10-second consistency threshold also satisfied the 1-second threshold. Note further that WinPMCC missed at least two GT set detections regardless of the employed threshold. Of course, the possibility still exists that false alarms were included in the GT set, partially because the detections are not categorized in terms of what types of events caused them. Event association, during which detections recorded by multiple sensor stations are associated to the same infrasound-producing event, is a function of network-level processing. This analysis, however, is limited to station-level processing.

Returning to the station-level ROC analysis in Figure 49, since the difference in the aforementioned trade-off is negligible for threshold consistencies between 1.0×10^{-6} sec and 0.01 sec, the use of thresholds within the following range is recommended: $0.01 \le c_n \le 1.0$. The false discovery rate (FDR) provides further insight into the burden on an analyst responsible for reviewing the list of detection families resulting from the use of such thresholds (pre-phase categorization and pre-repetitive source detection removal). Equation 24 quantifies the FDR for each of the thresholds analyzed, and the results are presented in Table 3. For the recommended threshold consistency range, an average of 81% of WinPMCC-declared detections are false alarms.

Table 3: FDR for each of the threshold consistencies c_n used to construct the pseudo-ROC curve in Figure 49.

Threshold Consistency (sec)	FDR
1.0×10^{-7}	N/A
1.0×10^{-6}	0.78
0.01	0.78
0.1	0.80
0.5	0.82
1.0	0.83
10	0.87

4.3 Optimum Family Size

To assist an analyst in his/her review of WinPMCC-declared detections, MAP and Bayes optimum thresholds are offered as decision guidelines for whether to include detections for further processing based upon the number of pixels comprising the families in question.

4.3.1 Maximum a Posteriori (MAP) Detection

Recall from Section 3.4 that the likelihood ratio test (LRT) in Eqn. 25 – shown here again –

$$\frac{p(z|T)}{p(z|R)} \overset{H_T}{\underset{H_R}{\gtrless}} \frac{P(R)}{P(T)} = \gamma \,, \tag{35}$$

relies, in part, on the conditional distributions of true event families $p(z|T)$ and non-event families $p(z|R)$. These conditional distributions are revisited here with the intention of more formally describing the lognormal and exponential distributions determined to have best fit the true event and non-event probability histograms respectively. Overlaid onto the true event probability histogram in Figure 30 is the

82

lognormal distribution

$$p(z|T) = f(z|\mu, \sigma) = \frac{1}{z\sigma\sqrt{2\pi}} e^{\frac{-(\ln z - \mu)^2}{2\sigma^2}}, \tag{36}$$

where $\mu = 2.94$ and $\sigma = 0.60$. This lognormal distribution describes the random

variable Z (number of pixels per family), whose logarithm is normally distributed given

that a true event has occurred. It was chosen to model the probability histogram because

it captures the histogram's positively skewed data. The lognormal mean μ is calculated

using

$$\mu = \ln\left(\frac{m^2}{\sqrt{var + m^2}}\right), \tag{37}$$

where m is the arithmetic mean, or first moment, of the GT data, and var is the variance.

For this GT data, $m = 22.61$ and $var = 221.92$. Likewise, the lognormal standard

deviation σ is calculated using [29]

$$\sigma = \sqrt{\ln\left(\frac{var}{m^2} + 1\right)}. \tag{38}$$

Opposite the true event family size distribution stands the non-event family

distribution. Overlaid onto its probability histogram in Figure 31 is the exponential

distribution

$$p(z|R) = f(z|\mu) = \frac{1}{\mu} e^{-\frac{z}{\mu}}, \tag{39}$$

where $\mu = 4.81$, or the average number of pixels comprising a non-event family. As

may intuitively be expected, the exponential distribution captures the higher likelihood of

smaller family sizes in the case of SOI absence. At least two pixels must be grouped

together before constituting a family. This minimum differs from the minimum number of pixels that were required to constitute a family while canvassing sensor data for detections to include in the GT set. Table 1 reveals that this minimum was set at 5 pixels. To soundly compare the true event and non-event family size distributions in the LRT, the minimum number of pixels required to constitute a family was therefore lowered to 2 for all time periods over which the GT set was constructed. No additional true detections existed at this adjusted family size minimum. Therefore, the true event and non-event distributions can now be compared.

Prior to determining the MAP family size threshold z_{t_MAP}, the remaining unknown terms in the LRT are the a priori probabilities of any WinPMCC-produced family corresponding to SOI presence $P(T)$ or SOI absence $P(R)$. Using the WinPMCC parameters and IDC-recommended 10-filter-per-decade configuration described in Table 1 as well as the default 0.1 sec consistency threshold, $P(T) = 0.37$ and $P(R) = 0.63$. The 63% a priori probability that any randomly chosen family belongs to the null hypothesis deviates from the 80% probability for the FDR in Table 3's consistency discussion for the following two reasons:

- MAP analysis occurs after phase-categorization has removed coherent noise detections caused by continuous/repetitive sources.

- The minimum number of grouped pixels required to constitute a family was lowered to 2 for the MAP analysis as opposed to 5 for the GT set.

This comparison can be made because *false alarm* categorizations within the ROC curve context are equivalent to this analysis's SOI *rejection* categorizations, as explained in

Figure 50: The maximum a posteriori threshold z_{t_MAP} based upon the ratio of true event $P(T) \cdot p(z|T)$ and non-event $P(R) \cdot p(z|R)$ likelihood functions in Eqn. 26. $P(T) = 0.37$ and $P(R) = 0.63$. Families of 12 or more pixels are more likely to indicate SOI presence than SOI absence.

Section 3.4's discussion of the different perspectives of the detection-declaring entities –

WinPMCC and analyst.

Ultimately, the a priori probabilities scale the true event and non-event

distributions accordingly, and the intersection of the resulting likelihood functions is

depicted in Figure 50. The MAP threshold for this examined data is 12 pixels per family.

Families of 12 or more pixels are more likely to indicate SOI presence (H_T) than SOI

absence (H_R), and the use of this threshold minimizes the number of H_T and H_R categorization errors.

4.3.2 Bayes Decision Criteria and Risk Minimization

If instead of minimizing the probability of categorization error, the goal is to minimize the cost-based *risk* associated with those categorization decisions, Bayes decision theory supplants the MAP approach. At the prerogative of a monitoring agency, such as the IDC, costs can be assigned to each of the following four potential events: detection, rejection, false alarm, and missed detection. Although the value for these costs is somewhat arbitrary, typically $C_{FA} > C_R$ and $C_{MD} > C_{TD}$. The subscripts in these inequalities refer to the cost of a false alarm, rejection, missed detection, and true detection respectively. For illustrative purposes, example costs may be assigned to a radar system whose purpose is to detect whether or not a missile has been launched as follows [30]:

- $C_R = 0$: no missile present, and correctly declare one not to be

- $C_{TD} = 10$: missile present, declare a missile to be present, and take action

- $C_{FA} = 20$: no missile present, but declare one to be

- $C_{MD} = 100$: missile present, but declare one not to be

The IDC can assign costs in a similar fashion, replacing "missile" in the above example with "SOI." The highest cost value should likewise be assigned to missing a detection.

Risk is modeled as a function of these costs and choices, and the minimization of that risk alters Eqn. 25's LRT as follows:

$$\frac{p(z|T)}{p(z|R)} \overset{H_T}{\underset{H_R}{\gtrless}} \frac{P(R)}{P(T)} \cdot \frac{C_{FA} - C_R}{C_{MD} - C_{TD}} . \tag{40}$$

The terms of Eqn. 40 are now rearranged so as to position the Bayes-scaled likelihood functions on either side of the inequality in the following manner [31]:

$$(C_{MD} - C_{TD}) \cdot P(T) \cdot p(z|T) \overset{H_T}{\underset{H_R}{\gtrless}} (C_{FA} - C_R) \cdot P(R) \cdot p(z|R) . \tag{41}$$

The true event and non-event family size likelihood functions are scaled by the previously assigned costs. As was the case for the MAP threshold, the optimum Bayes threshold z_{t_B} is marked by the graphical intersection of the functions on either side of Eqn. 41's inequality.

Given the higher cost assigned to missing a detection, the true event likelihood function $P(T) \cdot p(z|T)$ is scaled-up to a greater degree than the non-event likelihood function $P(R) \cdot p(z|R)$. Using the example costs assigned for the above radar system, the graphical intersection of these Bayes-scaled likelihood functions is shown in Figure 51. The rejection region refers to the area under the rejection likelihood function to the left of the threshold. The true detection region refers to the area under the true detection likelihood function to the right of the threshold. The false alarm region refers to the area under the rejection likelihood function to the right of the threshold. The missed detection region refers to the area under the true detection likelihood function to the left of the threshold. The resulting Bayes threshold $z_{t\_B} = 8$ implies basing H_T and

Figure 51: The optimum Bayes threshold z_{t_B} based upon the intersection of the true event $P(T) \cdot p(z|T)$ and non-event $P(R) \cdot p(z|R)$ likelihood functions in Eqn. 41 scaled by the example costs cited for the hypothetical radar system, which are $C_R = 0$, $C_{TD} = 10$, $C_{FA} = 20$, and $C_{MD} = 100$.

H_R categorization decisions on a fewer number of pixels per family than as suggested by

the MAP threshold z_{t_MAP}, which equals 12.

If instead the desire is to miss no more than 10% of detections, for instance, costs

can be assigned to allow no more than 10% of the area under the true event conditional

pdf $p(z|T)$ to fall to the left of the Bayes threshold. The costs assigned to the

hypothetical radar system coincidentally meet this stipulation. When the Bayes-scaled

likelihood functions intersect at 8 pixels, 7.65% of the area under $p(z|T)$ falls to the left

of the threshold, while intersection at 9 pixels violates the 10% stipulation. Specifically,

at $z_{t\_B} = 9$ pixels, 10.87% of the area under $p(z|T)$ falls to the left of the threshold.

4.3.3 Optimum Family Size Threshold Comparison

The trade-offs between threshold-based decisions and their true detection, missed

detection, false alarm, and rejection outcomes are presented for both the MAP and Bayes

approaches in the ROC analysis in Figure 52. ROC curves are constructed based both

upon the raw probability histogram data in Figure 30 and Figure 31 and the lognormal

and exponential conditional probability density functions that best fit the histograms'

data. Additionally, the probability of false alarm P_{FA} is converted to a FAR as a means to

offer insight into how many false alarms per day can be expected at various family size

decision thresholds.

Note that the ROC curves of Figure 52 cannot be compared on an "apples-to-

apples" basis with the pseudo-ROC curve of Figure 49 for the underlined reasoning in

Section 3.4. In addition to Section 3.4's delineation of detection-declaring entities

(WinPMCC *program* versus *analyst*), the pseudo-ROC curve is constructed by varying

the *consistency* threshold, whereas the optimum ROC curves are constructed by varying

the *family size* threshold. Moreover, the optimum family size ROCs analyze only those

families that remain *post*-repetitive source removal, whereas the consistency pseudo-

ROC analyzes all families *pre*-repetitive source removal. This distinction explains why

the optimum ROCs exhibit a lower FAR than the pseudo-ROC. The reason for the

distinction is due to the fact that the pseudo-ROC is meant to analyze the "detection

layer" in station-level processing – expounded upon in Section 3.1 and in Figure 24 –

Figure 52: ROC curves based upon the raw probability histogram data (red) in both Figure 30 and Figure 31 and the lognormal and exponential conditional probability density functions (blue) that best fit the histograms' data. The probability of detection P_D is plotted against both the probability of false alarm P_{FA} and an expected false alarm rate (FAR). Note that the ROC curve based upon the conditional PDFs is constructed by comparing the areas under these PDFs on either side of a varying family size threshold. The ROC curve based upon the raw probability histogram data reaches a maximum P_D of 0.936 to reflect the fact that WinPMCC misses at least 8 of the 125 GT set detections when the threshold consistency $= 0.1$ seconds.

prior to family "post-processing," of which repetitive source removal is a part. The

optimum ROCs, on the other hand, are meant to analyze the "detection layer" containing

family "post-processing." Of course, a pseudo-ROC curve can also be constructed for

the families that remain post-repetitive source removal. The *How To* guides in the

Appendix summarize the procedural steps, as outlined in this research, necessary to construct station-specific (pseudo) ROC curves and determine optimum family size decision thresholds.

For the *optimum MAP* threshold (12 pixels per family), the following decision outcomes apply: $P_D = 0.77$, $P_{MD} = 0.23$, $P_{FA} = 0.13$, and $P_R = 0.87$. $P_{MD} = 1 - P_D$ and therefore refers to the probability of missed detection. $P_R = 1 - P_{FA}$ and therefore refers to the probability of rejection. Relative to the costs mentioned earlier, the *optimum Bayes* threshold is $z_{t\_B} = 8$ pixels per family. Decisions based upon such a threshold imply the following decision outcomes: $P_D = 0.92$, $P_{MD} = 0.08$, $P_{FA} = 0.29$, and $P_R = 0.71$. The Bayes approach seeks to minimize the expected value of the decision outcome costs. Since the highest cost was assigned to a missed detection, the least likely outcome for Bayes threshold-based decisions is indeed a missed detection.

4.4 SNR Stress Tests and Detector Failure

WinPMCC is tested in decreasing SNR conditions to determine the absolute limit of its detection capability. This limit, as determined for synthetic data and synthetic arrays as well as for real data and existing arrays, is defined as the post-filtered SNR at which WinPMCC fails to detect a SOI at least 90% of the time or, equivalently, when $P_D \leq 10\%$. P_D is based upon the use of 20 different noise realizations at each tested SNR. The failure level therefore refers to the SNR at which WinPMCC misses the SOI for at least 18 of the 20 different noise realizations. Recall from Sections 2.5 and 3.5 that PMCC calculations on a time window of data do not commence until after that waveform data is filtered according to the filter configuration established in WinPMCC's "Window

and Filter Parameter" settings. The filtering operation of "Filter 16," as discussed in Sections 3.5.1 and 3.5.3, is the most appropriate to duplicate in order to determine the post-filtered SNR for both the synthetic and real data stress tests. The motivation behind such work stems from the blinding effect of high-amplitude incoherent noise that wind bursts introduce for the infrasound stations located in particularly windy global environments.

4.4.1 Synthetic Data SNR Stress Test

The synthetic SNR stress tests rely on WinPMCC's ability to detect the Pierce Blast in Figure 32 amidst additive white Gaussian noise (AWGN), modeled with MATLAB's *randn* function, on the 9-element synthetic array in Figure 43. For this configuration, WinPMCC fails to detect the SOI at least 90% of the time when the post-filtered SNR drops below 3 dB. To ascertain whether deterioration in detection capability accelerates for configurations with fewer elements, sensors are removed from potentially participating in the PMCC detection algorithm. Synthetic configurations that resemble the configurations of the 5 stations whose recorded detections were used to build the GT set are created from sub-arrays of the 9-element synthetic array. These five real arrays and their synthetic look-alikes are pictured in Figure 45, and the SNR stress test results for the look-alikes are presented in Table 4.

As Figure 20 revealed, stations with fewer sensor elements struggle to as accurately *estimate* the angle of arrival for SOIs traveling as plane waves. The progressive inclusion of more sensors in the PMCC detection algorithm was also introduced in Section 2.3 in the context of improving SOI azimuth and velocity *estimates*. These SNR stress test results, however, suggest only a marginal advantage for arrays with

92

more elements in terms of the ability to *detect* SOIs in environments with high-amplitude incoherent noise.

4.4.2 Real Data SNR Stress Test

In much the same way as the synthetic test, the real data SNR stress test proceeds with determining WinPMCC's ability to detect the SOI in Figure 46 amidst sensor-recorded noise on the 5-element BRD array. A one-minute excerpt of the original data recorded by sensor BRD00 – the top element in Figure 47 – is pictured in Figure 53. The last 15 seconds of this excerpt is the time window in which the detection lies. The data are shown pre and post-filtered. The array-averaged post-filtered SNR, as approximated by the STA/LTA SNR proxy in Eqn. 34, is 20 dB.

In the manner outlined by Section 3.5.3, 60 seconds of pre-filtered noise from a detectionless window of BRD00 data is overlaid onto the original pre-filtered data in Figure 53. This process is repeated for the other four sensors in the array, and the amplitude of that overlaid noise is increased until WinPMCC fails to detect the SOI. Instead of overlaying different *randn* noise realizations as in the synthetic stress test, different detectionless windows of data are overlaid onto the SOI and amplified until

Table 4: Synthetic SNR Stress Test Results. WinPMCC fails for post-filtered SNRs lower than what is indicated at a 90% rate. "Filter 16" with the response shown in Figure 34 is used to determine post-filtered SNRs.

Sensor Array	Number of Array Elements	Pre-Filtered Failure SNR (dB)	Post-Filtered Failure SNR (dB)
Full Synthetic	9	-9	3
Synthetic BRD	5	-8	4
Synthetic CHN	4	-8	4
Synthetic KSG	4	-8	4
Synthetic I30	6	-9	3
Synthetic I45	3	-8	4

Figure 53: Pre-filtered (top) raw sensor data recorded by sensor element BRD00 on the BRD array. The SOI is located between 45 – 60 seconds, as is apparent on the post-filtered (bottom) waveform. "Filter 16" (shown in Figure 34) from WinPMCC's 20-filter configuration illustrated in Figure 9 maximizes the post-filtered SNR, which the STA/LTA approximation indicates is 20 dB.

WinPMCC misses the detection for at least 90% of attempts. The post-filtered SNR at which this 90% missed detection rate occurs is 2 dB. Figure 54 illustrates example pre and post-filtered sensor data at this BRD failure SNR.

4.5 Summary and Impact of Results

The pseudo-ROC and optimum threshold procedures outlined in Chapters III and IV can be applied to individual stations at the discretion of the IDC. GT sets built from detections on particular stations may be less prone to the inclusion of non-event families due to the IDC's access to event-confirmed detections at the output of network-level

94

processing. These research efforts attempted to compensate for the unavailability of

network-level processing by requiring that GT set detections be confirmed by at least two

of the three detection methods employed. With a completed GT set, the IDC (or other

interested monitoring agency, such as the Air Force Technical Applications Center

(AFTAC)) can repeat the procedures as described in this document to determine the

consistency dependent trade-off between P_D and FAR as well as optimum decision

thresholds based upon the true event $p(z|T)$ and non-event $p(z|R)$ conditional

Figure 54: Pre-filtered (top) sensor data consisting of 60 seconds of noise from a detectionless window of data overlaid on top of the raw data in Figure 53. Although only sensor element BRD00 is shown, this same process of overlaying noise from detectionless windows of data is repeated on BRD's other 4 sensor elements. The amplitude of the overlaid noise is increased until WinPMCC fails to detect the SOI, located between 45 – 60 seconds. "Filter 16" (shown in Figure 34) from WinPMCC's 20-filter configuration illustrated in Figure 9 maximizes the post-filtered SNR, which the STA/LTA approximation indicates is 2 dB. For post-filtered (bottom) waveforms with SNRs below this 2 dB failure level, WinPMCC fails to detect the SOI at greater than a 90% rate.

probability density functions.

The IDC can then judge array performance by comparing these station-specific curves and thresholds. Steeper pseudo-ROC curves imply better performance, as does less overlap between the (scaled) true event and non-event conditional probability density functions. With less overlap, the decision threshold, whether it is based upon MAP or Bayes criteria, is less susceptible to false alarm and missed detection categorization errors.

Monitoring agencies can further assess individual array geometries and station performance by evaluating WinPMCC's ability to detect SOIs in progressively noisier environments. Noisy environments can either be simulated synthetically or with real data in the manner outlined by Chapters III and IV. The results for such methods may not coincide exactly, as they did not – 2 dB difference in post-filtered failure SNR level – for the fielded BRD array and its synthetic look-alike. The reasons for synthetic and real data result mismatches are varied. First, the STA/LTA approximation for SNR is just that, an *approximation*. Moreover, noise for the synthetic SNR stress tests is *modeled* with MATLAB's *randn* function, and models are imperfect representations of reality. For instance, *randn* noise vector realizations are meant to simulate AWGN when, in fact, the noise profiles of specific arrays may actually exhibit *colored* noise. Further, the synthetic arrays in Figure 45 have smaller apertures than their real counterparts. In addition, there are amplitude and frequency content differences between the synthetic Pierce Blast SOI and the real data SOI waveforms. Finally, synthetic SOI propagation is modeled as a perfect plane wave, whereas real SOIs may deviate from plane wave

propagation due to the path-altering effects underscored in Section 2.8. Nevertheless, both the synthetic and real data SNR stress test methods provide valuable insight into WinPMCC's detector limitations in noisy environments.

V. Conclusions and Recommendations

This chapter begins by restating the quantitative results for the consistency threshold (c_n) Receiver Operating Characteristic (ROC) analysis, the proposed maximum a posteriori (MAP) and Bayes optimum family size thresholds, and the signal-to-noise ratio (SNR) stress tests. These results are then discussed within the context of how they contribute to the missions of the International Data Centre (IDC), which currently employs the Progressive Multi-Channel Correlation (PMCC) algorithm in an operational setting, and the Air Force Technical Applications Center (AFTAC), which is exploring integrating PMCC into its detection repertoire. Finally, recommendations for how future work can build upon this research's performance evaluation procedures are suggested.

5.1 Results Summary

- Pseudo-ROC curve analysis explored the consistency-dependent trade-off between the probability of detection (P_D) and the false alarm rate (FAR) for WinPMCC-produced families consisting of at least 5 elementary detections (pixels). Pixel creation is described in detail in Figure 6. The results suggested that consistency thresholds above 1.0 second increase FAR without any appreciable increase in P_D, thus unnecessarily increasing the burden on an analyst reviewing the list of WinPMCC-produced detections. The false discovery rate (FDR), or the metric quantifying the percentage of program-produced detections that are false alarms, peaked at 87% for a 10 second consistency threshold. On the other hand, thresholds between 1.0×10^{-6} and 0.01 seconds exhibited no apparent difference in the aforementioned trade-off, whereas thresholds increasingly below 1.0×10^{-6} seconds

caused WinPMCC's P_D to rapidly decline. Specifically, the use of a 1.0×10^{-6}

second threshold implied a 20% missed detection rate, and a 1.0×10^{-7} second

threshold produced no detections at all. Given these results, thresholds within the

following range are recommended for elementary pixel detection: $0.01 \leq c_n \leq 1.0$.

- For an analyst responsible for making decisions on which pixel *families* to preserve

  for network-level processing, MAP and Bayes *optimum* decision thresholds based

  upon the sizes of the detection families in question are proposed. For a 0.1 second

  consistency threshold and the WinPMCC settings expressed in Table 1, the *MAP*

  decision threshold is *12 pixels*. Any family comprised of at least 12 pixels is more

  likely to indicate signal-of-interest (SOI) presence than absence (and vice versa for

  families comprised of fewer than 12 pixels). Thus, the MAP threshold minimizes the

  probability of error, defined as $P_{error} = P_{FA} + P_{MD}$. P_{FA} refers to the probability of

  false alarm, and P_{MD} refers to the probability of missed detection. In Bayes

  optimization, unitless costs are assigned to all possible decision outcomes – true

  detection, rejection, false alarm, and missed detection. The Bayes optimum threshold

  minimizes decision *risk*, defined as the expected value of the assigned costs. For true

  detection, rejection, false alarm, and missed detection costs of 10, 0, 20, and 100

  respectively, the *Bayes optimum threshold* is 8 *pixels*. When families comprised of 8

  or more pixels are categorized as true event families and those with fewer than 8

  pixels as non-event families, $P_{MD} < 8\%$ and $P_{FA} = 29\%$, which equates to about 30

  false alarms per day. Referencing Figure 49 for an apples-to-apples error comparison

  with the consistency-dependent pseudo-ROC analysis, use of the same default

consistency (0.1 sec) and a minimum family size of 5 pixels, $P_{MD} = 5\%$ with an associated FAR of about 240 false alarms per day.

- The SNR stress tests provided insight into the conditions under which high-amplitude incoherent noise renders sensor arrays blind to potential SOIs. Array "blindness" is defined as the post-filtered SNR at which WinPMCC's $P_{MD} \geq 90\%$. For the 9-element synthetic array and the synthetic BRD (5 elements), CHN (4 elements), KSG (4 elements), I30 (6 elements), and I45 (3 elements) array look-alikes, the missed detection rate exceeded 90% when the post-filtered SNR dropped below 3, 4, 4, 4, 3, and 4 dB respectively. To compare the synthetic stress test results with those for real data recorded by the operational BRD array, the missed detection rate was determined to exceed 90% when the post-filtered SNR dropped below 2 dB. These results suggest that arrays with more elements have only a marginal advantage in terms of PMCC's ability to *detect* SOIs in high-amplitude noise environments. Note that these stress tests did not measure PMCC's ability to accurately *estimate* a SOI's propagating velocity or azimuth. Previous work has demonstrated that signal parameter estimation generally improves for arrays with more elements, but the level of this improvement has not yet been quantified in deteriorating SNR conditions.

5.2 Research Contributions

Monitoring agencies that maintain infrasound stations, such as AFTAC or the IDC, can use the performance evaluation procedures established in this document to assist in the performance improvement of individual stations. The IDC, for instance, can build reliable, station-specific ground truth (GT) sets by assembling the event-confirmed

detections at the output of network-level processing. These research efforts attempted to compensate for the unavailability of network-level processing by requiring that GT set detections be confirmed by at least two of the three detection methods (WinPMCC, InfraMonitor, FK Trend) employed, an obstacle the IDC need not overcome. The IDC can then construct station-specific pseudo-ROC curves by following the procedures outlined in Chapter III, summarized in the *How To* guide in the Appendix. The performance of different array geometries can be judged by comparing their individual pseudo-ROC curves. Steeper curves, wherein the trade-off between P_D and FAR is more favorable, imply better performance. The assumption here is that, in order to solely compare the "steepness" of station-specific pseudo-ROC curves, the FARs for each must be equivalent. For example, identical curves can still signify performance disparity if each curve is plotted against respective FARs that are unique to a certain geometry, location, and/or season. In practice, therefore, pseudo-ROC steepness will have to be considered jointly with relative FARs when judging array geometry performance. For a specific consistency threshold, the expected burden on an analyst responsible for reviewing the list of WinPMCC-produced detections can further be quantified with the false discovery rate (FDR).

Recall that analysts review *families* of detections, not merely the consistency-satisfied elementary detection *pixels*. The IDC does not currently use any family-characterizing statistic as a detection threshold, but this research's results suggest that *family size* is a viable candidate. The purpose of such a threshold is twofold. First, the onerous, time-consuming process of reviewing detections becomes streamlined when a detection threshold provides guidance as to the *likelihood* (MAP approach) or *risk* (Bayes

cost criteria) of declaring whether families indicate SOI presence or SOI absence. For reference, the optimum threshold analysis culminated in the ROC curve comparison in Figure 52, revealing the true detection, rejection, false alarm, and missed detection rates for each approach. Second – and perhaps more insightful than the previously described consistency-based station evaluation method – individual station performance can alternately be evaluated by comparing the station-specific overlap of the MAP or Bayes-scaled density functions "fit" to the true event and non-event probability histogram families. Less overlap reduces the minimum achievable categorization error sum of P_{FA} and P_{MD}, P_{error}, in the case of the MAP approach. Likewise, in the case of the Bayes approach, less overlap reduces the minimum achievable risk when the highest costs are assigned to missed detection and false alarm decision outcomes, as they often are in practice. The *How To* guide in the Appendix also summarizes the procedural steps necessary to determine station-specific optimum thresholds. The International Monitoring System (IMS) can ultimately improve the efficacy of its infrasound network by updating the array geometries of its infrasound network's stations to reflect the best performing of the tested geometries.

The final method discussed by which array geometries can be evaluated was via SNR stress tests. The results, however, did not suggest a clear advantage for any single array configuration in deteriorating SNR conditions. They did, however, provide previously unknown insight into PMCC's SNR detection capability limitations.

5.3 Recommendations for Future Work

As there are a multitude of WinPMCC parameter settings, these evaluation methods can be repeated under different initial program settings to determine setting-dependent performance disparity. For instance, multiple pseudo-ROC curves can be drawn where each one reflects a different filter configuration, thereby providing a method to compare detector performance with different filtering schemes.

The GT set detections, upon which the majority of this research relies, can perhaps be divided into low, medium, and high F-statistic or SNR subsets. From there, the MAP and Bayes optimum family size detection thresholds can be determined for each of these subsets, thereby revealing the degree to which they may or may not deviate. Moreover, optimum decision thresholds based upon family-characterizing statistics other than *size* can be independently determined. For example, true event and non-event probability histograms can be created by sorting GT set and noise set detections according to their F-statistics rather than their family sizes. The reliability of analysts' decisions as to whether families should be preserved for network-level processing can only benefit from access to multiple decision thresholds based upon various family attributes.

From a detector limitation perspective, the real data SNR stress tests can be performed on the remaining untested fielded arrays (only the BRD array was stress tested). Further, a more direct comparison between the synthetic and real data SNR stress tests can ensue if synthetic arrays are designed so that their apertures are more in line with the apertures of the fielded arrays studied. Such efforts also hint at the possibility of testing not only PMCC performance for various array geometries, but also for various

array apertures. Quantifying PMCC's ability to accurately estimate the azimuth and velocity of a propagating infrasound signal on these array structures in deteriorating SNR conditions could be another avenue in which to investigate. Finally, the definition of array "blindness" could be adjusted to lower P_{MD} levels to reveal whether any single configuration has a clear detection advantage at these other levels.

Appendix

How to Create a Consistency-Dependent Station-Specific Pseudo-ROC Curve

1) <u>Decide which consistency thresholds (c_n) to test</u>. Those chosen for the analysis presented in this document are specified in Figure 49.

2) <u>Build a ground truth (GT) set of event arrivals on the infrasound station for which the pseudo-ROC curve will be constructed</u>. To achieve a representative sample of arrivals, it is desirable that the GT set consist of at least 100 confirmed detections. If access to the International Data Centre's (IDC) Standard Event Lists (SELs) is available, make note of events in the proximity of the station to be tested. If access to either SELs or a specific monitoring agency's list of events (such as AFTAC's) is unavailable, follow the process of building the GT set as outlined in Section 3.2. Regardless of how the GT set is constructed, do the following:

   a. Fix WinPMCC settings that will *not* be varied throughout the ROC-building process, such as the filter parameters, detection parameters (other than c_n), and families settings in Table 1. These settings will be used throughout the creation of one pseudo-ROC curve.

   b. Keeping in mind that this pseudo-ROC will be evaluating the *consistency-dependent* trade-off between the probability of detection (P_D) and the false alarm rate (FAR), run the WinPMCC program at the *highest* (most lenient) c_n that will be tested.

   c. Arrival WinPMCC detections on the station that can be associated with either the SEL events or AFTAC-confirmed events can be added to the GT set.

3) <u>Note detection characteristics for each of the GT event detections recorded by the station</u>. WinPMCC outputs a number of characteristics, such as family size, F-stat, and the number of sensor elements within the array participating in the detection. Be sure to note the time of arrival and the window length of the produced family (ex. 15 second detection). These values may prove useful in identifying GT set detections throughout the process of plotting ROC curve points. It is recommended that the detection-characterizing data be organized in a matrix within Excel. Excel is suggested because, if needed, MATLAB can import and analyze an Excel spreadsheet.

4) <u>Decide on a standard amount of time for which to run WinPMCC before and after a known GT set detection</u>. For example, if a GT set detection occurs at noon on a particular day, decide to run WinPMCC from 15 min prior to the arrival through 15 min after, i.e. from 11:45-12:15. This stipulation is offered to ensure WinPMCC canvasses time windows of data within which there should be no detections, thus allowing a FAR to be determined based upon instances in which WinPMCC flags a time window of data as a detection, but should not.

5) <u>Run WinPMCC over all time windows established in the previous step at each of the c_n's established in Step 1</u>. For each c_n, note the number of false alarms* as well as the number of successfully detected GT set detections and the number of unsuccessfully detected, or "missed," detections. Missed detections refer to instances in which WinPMCC failed to detect a GT infrasound arrival. Now if, for example, seven c_n were analyzed, there will be seven pseudo-ROC points composing the curve.

Each point consists of an overall P_D and a total false alarm number based upon the use of an individual c_n.

6) Decide what *rate* by which to quantify the "total false alarm number." Recall from Eqn. 23 that the analysis presented in this document chose to quantify the FAR on a *per day* basis.

7) Graph the points on a P_D versus FAR plot, thereby creating a pseudo-ROC curve, as in Figure 49.

The following approaches can be taken from here:

- Repeat this step-by-step procedure for another station(s) and compare the resulting station-specific pseudo-ROC curves to judge the performance disparity of station array configurations/geometries.

- Alter one of the previously fixed settings in Step 2a above, such as filter configuration, and create another pseudo-ROC curve for the same station, varying c_n's in the same manner. WinPMCC's relative performance based upon how individual settings are tuned can now be judged by overlaying and comparing the original and newly created pseudo-ROC curves.

Note in Step 5 that families produced by repetitive noise sources, such as ocean swell microbaroms, are removed from the WinPMCC detection list in time-frequency space in the manner described in Section 2.5 and presented visually in Figure 14. The analysis presented within this document chose to count false alarms prior to this detection list "cleaning." However, since coherent noise detections produced by repetitive sources

107

would be eliminated anyway, it makes sense to only count false alarms that remain after detection "cleaning."

How to Determine Station-Specific Optimum Family Size Thresholds

1) <u>Build a GT set of event arrivals on the infrasound station for which the optimum family size threshold(s) will be determined</u>. To achieve a representative sample of arrivals, it is desirable that the GT set consist of at least 100 confirmed detections. If access to the International Data Centre's (IDC) Standard Event Lists (SELs) is available, make note of events in the proximity of the station to be tested. If access to either SELs or a specific monitoring agency's list of events (such as AFTAC's) is unavailable, follow the process of building the GT set as outlined in Section 3.2. Regardless of how the GT set is constructed, do the following:

   a. Fix *all* WinPMCC settings, such as the filter parameters, detection parameters (including c_n), and families settings in Table 1. These settings will be used throughout the process of determining the optimum family size threshold(s).

   b. Arrival WinPMCC detections on the station that can be associated with either the SEL events or AFTAC-confirmed events can be added to the GT set.

2) <u>Note detection characteristics for each of the GT event detections recorded by the station</u>. WinPMCC outputs a number of characteristics, such as family size, F-stat, and the number of sensor elements within the array participating in the detection. Be sure to note the time of arrival and the window length of the produced family (ex. 15 second detection). These values may prove useful in identifying GT set detections throughout the process of determining the optimum threshold(s). It is recommended

that the detection-characterizing data be organized in a matrix within Excel. Excel is

suggested because, if needed, MATLAB can import and analyze an Excel

spreadsheet.

3) <u>Decide on a standard amount of time for which to run WinPMCC before and after a</u>

<u>known GT set detection</u>. For example, if a GT set detection occurs at noon on a

particular day, decide to run WinPMCC from 15 min prior to the arrival through

15 min after, i.e. from 11:45-12:15. This stipulation is offered to ensure WinPMCC

canvasses time windows of data within which there should be no detections.

Instances in which WinPMCC flags a time window of data as a detection (by

producing a family) that cannot be associated with a GT set arrival are rejections, or

non-event families.

4) <u>Run WinPMCC at the decided upon settings from Step 1a over all time windows</u>

<u>established in Step 3</u>. There is now a finite set of WinPMCC-produced detections, or

families, within these time windows.

5) <u>Sort these detections into either GT set true event detections or (rejection) non-event</u>

<u>detections</u>. *Prior to sorting*, eliminate obvious detections produced by repetitive

noise sources, for this is accomplished anyway during *post-PMCC processing*.

Detection list cleaning was discussed in greater detail in Section 2.5 and is presented

visually in Figure 14.

6) <u>Note the family sizes, or the number of pixels per family, for each of the GT set true</u>

<u>event detections</u>. <u>Do the same for the non-event families</u>.

7) Plot probability histograms, also known as probability mass functions, that show the distribution of family sizes for both the GT set true event detections and non-event detections, as in Figure 30 and Figure 31.

8) Decide which probability density functions (pdf's) best characterize the data presented by the probability histograms. For the analysis presented in this document, the lognormal pdf in Eqn. 36 best characterizes the true event data, and the exponential pdf in Eqn. 39 best characterizes the non-event data.

9) Given the finite list of families established in Step 4, determine the probabilities of any randomly chosen family belonging to either the true event or non-event sets. These are the "a priori" probabilities described in the likelihood ratio test (LRT) of Eqn. 25.

10) Scale the true event conditional pdf established in Step 8 by the a priori probability that any randomly chosen family belongs to the GT set true event detection list. Likewise, scale the non-event conditional pdf by the a priori probability that any randomly chosen family belongs to the non-event detection list.* "Scaling" means to multiply a conditional pdf at each of its sample points by the appropriate a priori probability.

11) Graph the scaled true event and non-event conditional pdfs, also known as likelihood functions, on the same plot, as in Figure 50. Note the intersection, which should be rounded up to the nearest pixel. This intersection marks the *maximum a posteriori* (MAP) threshold, which minimizes the probability of categorization error, P_{error}, defined as $P_{error} = P_{FA} + P_{MD}$. P_{FA} refers to the probability of false alarm, and P_{MD} refers to the probability of missed detection. Families that are comprised of at least as

many pixels as the MAP threshold are more likely to be true event families than non-event families. Families comprised of fewer pixels than the MAP threshold are more likely to be non-event families.

Following are some additional thoughts:

- Optimum decision thresholds based upon family-characterizing statistics other than *size* can be independently determined. For example, true event and non-event probability histograms can be created by sorting GT set and noise set detections according to their F-statistics rather than their family sizes. MAP and Bayes optimum F-statistic thresholds can be determined in the same manner as the optimum family size thresholds were determined.

- The reliability of analysts' decisions as to whether families should be preserved for network-level processing can only benefit from access to multiple decision thresholds based upon various family attributes.

Note in Step 10 that if the conditional pdf's are also scaled by Bayes costs, as in Eqn. 41, their intersection denotes the optimum Bayes threshold. When categorization decisions are made based upon this threshold, decision risk – defined as the expected value of the Bayes costs – is minimized.

Bibliography

[1]   University of Western Ontario, Department of Physics and Astronomy, "Meteor
       Infrasound," [Online]. Available:
       http://meteor.uwo.ca/research/infrasound/is_meteorIS.html. [Accessed 30
       January 2013].

[2]   A. Le Pichon, E. Blanc, and A. Hauchecorne, Eds., Infrasound Monitoring for
       Atmospheric Studies, New York: Springer, 2010.

[3]   Commissioner for Atomic Energy (CEA), *WinPMCC version 4.0c3,* Department of
       Analysis, Monitoring, Environment (DASE), 2010.

[4]   "Preparatory Commission for the Comprehensive Nulcear-Test-Ban Treaty
       Organization," [Online]. Available: http://www.ctbto.org/specials/who-we-are/.
       [Accessed 4 January 2013].

[5]   D. Clauter, *Memo Regarding PMCC Evaluation Effort,* Air Force Technical
       Applications Center, 2012.

[6]   M. A. Garcés and C. H. Hetzer, "Optimizing the Progressive Multi-Channel
       Correlation Detector for the Discrimination of Infrasonic Sources," in *25th
       Seismic Research Review - Nuclear Explosion Monitoring: Building the
       Knowledge Base*, Tuscon, 2003.

[7]   "Department Assessment and Monitoring Environment: MB2000 and MB2005
       Microbarometers," [Online]. Available: http://www-
       dase.cea.fr/public/dossiers_thematiques/microbarometres/description_en.html.
       [Accessed 13 January 2012].

[8]   A. Le Pichon, A. Vergoz, E. Blanc, J. Guilbert, L. Ceranna, L. Evers and N. Brachet,
       "Assessing the Performance of the International Monitoring System Infrasound
       Network: Geographical Coverage and Temporal Variabilities," *Journal of
       Geophysical Research,* vol. 114, 2009.

[9]   D. N. Green and D. Bowers, "Estimating the Detection Capability of the
       International Monitoring System Infrasound Network," *Journal of Geophysical*

Research, vol. 115, 2010.

[10] A. Le Pichon and Y. Cansi, "Progressive Multi-Channel Correlation: Technical Documentation, CTBTO 2003-0269/POGGIO," CTBTO, 2003.

[11] Y. Cansi and Y. Klinger, "An Automated Data Processing Method for Mini-Arrays," European-Mediterranean Seismological Centre, 1997.

[12] Y. Cansi and A. Le Pichon, "Infrasound Event Detection Using the Progressive Multi-Channel Correlation Algorithm," in *Handbook of Signal Processing in Acoustics*, New York, Springer, 2008.

[13] R. H. Shumway, S.-E. Kim and R. R. Blandford, "Nonlinear Estimation for Time Series Observed on Arrays," in *Asymptotics, Nonparametrics, and Time Series*, New York, Marcel Dekker, 1999, pp. 227-258.

[14] R. Kemerait, Interviewee, *Discussion with Dr. Kemerait Regarding the History of the IMS and IDC Data Processing.* [Interview]. 28 June 2012.

[15] A. Le Pichon, R. Matoza, N. Brachet and Y. Cansi, "Recent Enhancements of the PMCC Infrasound Signal Detector," *InfraMatics,* pp. 5-8, September 2010.

[16] A. Le Pichon and Y. Cansi, "PMCC for Infrasound Data Processing," *InfraMatics,* June 2003.

[17] R. H. Shumway and S. Kim, "Signal Detection and Estimation of Directional Parameters for Multiple Arrays," University of California, Davis, 2001.

[18] C. Szuberla and J. Olson, "Uncertainties Associated with Parameter Estimation in Atmospheric Infrasound Arrays," *Journal of the Acoustical Society of America,* vol. 115, pp. 253-258, 2004.

[19] J. Olson, "Least-Squares Estimation of the Azimuth and Velocity of Plane Waves Using Infrasound Array Data," Geophysical Institute, University of Alaska, 2004.

[20] R. Shumway, "Advances in Mixed Signal Processing for Regional and Teleseismic Arrays," University of California, Davis, 2006.

[21] R. Shumway, E. Smart and D. Clauter, "Mixed Signal Processing for Regional and Teleseismic Arrays," *Bulletin of the Seismological Society of America,* vol. 98, pp. 36-51, 2008.

[22] J. Olson, "Infrasound Signal Detection using the Fisher F-Statistic," *InfraMatics,* pp. 1-7, June 2004.

[23] D. Brown, R. Whitaker, B. Kennett and C. Tarlowski, "Automatic Infrasonic Signal Detection Using the Hough Transform," *Journal of Geophysical Research,* vol. 113, 2008.

[24] S. Arrowsmith, R. Whitaker, S. Taylor, R. Burlaeu, B. Stump, M. Hedlin, G. Randall, C. Hayward and D. ReVelle, "Regional Monitoring of Infrasound Events Using Multiple Arrays: Application to Utah and Washington State," *Geophysical Journal International,* vol. 175, pp. 291-300, 2008.

[25] Air Force Technical Applications Center (AFTAC), *SeaTools,* Patrick Air Force Base, Version 3.0.91.

[26] T. Fawcett, "ROC Graphs: Notes and Practical Considerations for Researchers," Kluwer, Palo Alto, 2004.

[27] B. Sklar and P. K. Ray, Digital Communications: Fundamentals and Applications, New York: Pearson, 2009.

[28] M. Temple, *Digital Communications I Lecture Notes,* Dayton: Air Force Institute of Technology, 2010.

[29] MATLAB, *version 7.14.0,* Natick: Mathworks, Inc., 2012.

[30] R. Martin, *EENG 663: Signal Detection and Estimation,* Air Force Institute of Technology, 2012.

[31] S. M. Kay, Fundamentals of Statistical Signal Processing: Volume II, Detection Theory, Upper Saddle River: Prentice Hall, 1998.

Vita

Anthony Runco ████████████████████████████. He graduated as the

Valedictorian from Chaminade High School and subsequently attended Yale University

where he majored in electrical engineering. He spent many of his high school and

college summers working as an ocean lifeguard for the Village of Atlantic Beach, his

hometown. At Yale, he co-founded a fitness society and participated in the Air Force

Reserve Officer Training Corps (AFROTC). He graduated *cum laude* with a B.S. degree

in 2011 with further *Distinction in the Major* honors. He received Yale's Edward O.

Lanphier Memorial Prize, which is awarded to the top graduating electrical engineer.

Prior to commissioning, he held positions at both Northrop Grumman and J.P. Morgan.

He will graduate from the Air Force Institute of Technology (AFIT) in March 2013 with

a Master's in electrical engineering. During his time at AFIT, he served as the Vice

President for AFIT's chapter of Eta Kappa Nu (HKN), the national honor society for

electrical and computer engineers. His next Air Force assignment involves working at

the Space Superiority Systems Directorate in Los Angeles, CA.

REPORT DOCUMENTATION PAGE				Form Approved OMB No. 074-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of the collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) 21-03-2013	2. REPORT TYPE Master's Thesis	3. DATES COVERED (From – To) Sept 2011 – March 2013

4. TITLE AND SUBTITLE Detection Optimization of the Progressive Multi-Channel Correlation Algorithm Used in Infrasound Nuclear Treaty Monitoring	5a. CONTRACT NUMBER
	5b. GRANT NUMBER
	5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S) Runco, Anthony M., Second Lieutenant, USAF	5d. PROJECT NUMBER
	5e. TASK NUMBER
	5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way, Building 640 WPAFB OH 45433	8. PERFORMING ORGANIZATION REPORT NUMBER AFIT-ENG-13-M-42

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Technical Applications Center (AFTAC/TTR) Dr. Dean Clauter 1030 S Hwy A1A, Building 989 Patrick AFB, FL 32925 Phone: (321) 494-5266 E-mail: dean.clauter@us.af.mil	10. SPONSOR/MONITOR'S ACRONYM(S) AFTAC/TTR
	11. SPONSOR/MONITOR'S REPORT NUMBER(S)

DISTRIBUTION STATEMENT A. APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

13. SUPPLEMENTARY NOTES
This material is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

14. ABSTRACT
This thesis develops methods to determine optimum detection thresholds for the Progressive Multi-Channel Correlation (PMCC) algorithm used by the International Data Centre (IDC) to perform infrasound station-level nuclear-event detection. Receiver Operating Characteristic (ROC) curve analysis is used with real ground truth data to determine the trade-off between the probability of detection (P_D) and the false alarm rate (FAR) at various consistency detection thresholds. Further, statistical detection theory via maximum a posteriori and Bayes cost approaches is used to determine station-level optimum "family" size thresholds of grouped detection "pixels" with similar signal attributes (i.e. trace velocity, azimuth, time of arrival, and frequency content) before the detection should be considered for network-level processing. Optimum family sizes are determined based upon the consistency threshold and filter configuration used to filter sensor data prior to running the detection algorithm. Finally, this research generates synthetic signals for particular array configurations, adjusts the signal-to-noise ratio (SNR) to determine the SNR failure levels for the PMCC detection algorithm, and compares this performance to the performance of fielded infrasound stations with similar configurations.

15. SUBJECT TERMS
Progressive Multi-Channel Correlation; PMCC; infrasound; nuclear treaty monitoring; detection; optimization

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Louthain, James A., Lt Col, Ph.D,USAF
a. REPORT	b. ABSTRACT	c. THIS PAGE	UU	130	19b. TELEPHONE NUMBER (Include area code) (937) 255-3636, x 4620 (james.louthain@afit.edu)
U	U	U			

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39-18